

To BeOS Or Not

Java Trees

Black-Holing Spam

PERFORMANCE COMPUTING

THE MAGAZINE FOR THE NEW DATA CENTER

Managing Intranets

**The Outsourced
Intranet—
Minding Your
ISPs & Qs**

**RS/6000 HA-H70:
IBM's Cluster
With Muster**

**Red Hat's
Latest Linux**

November 1999
US \$3.95 Canada \$4.95
www.performance-computing.com

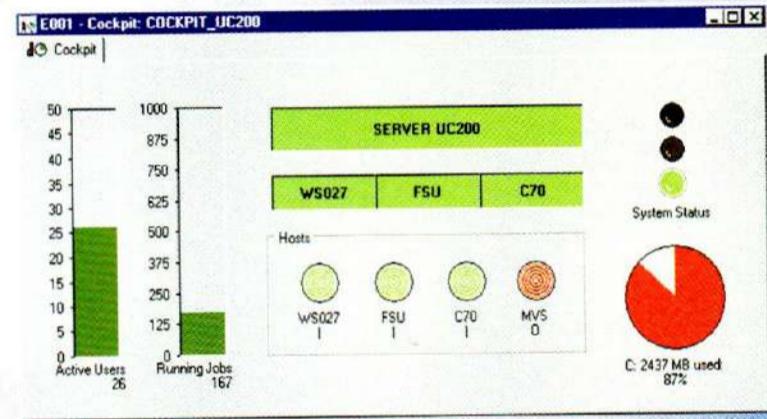
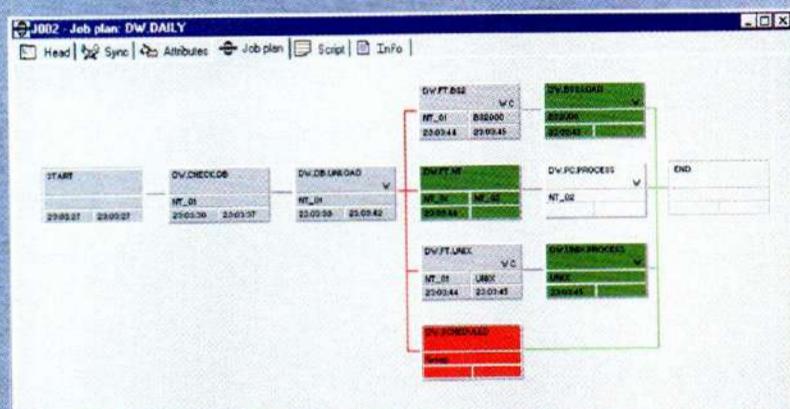


11>

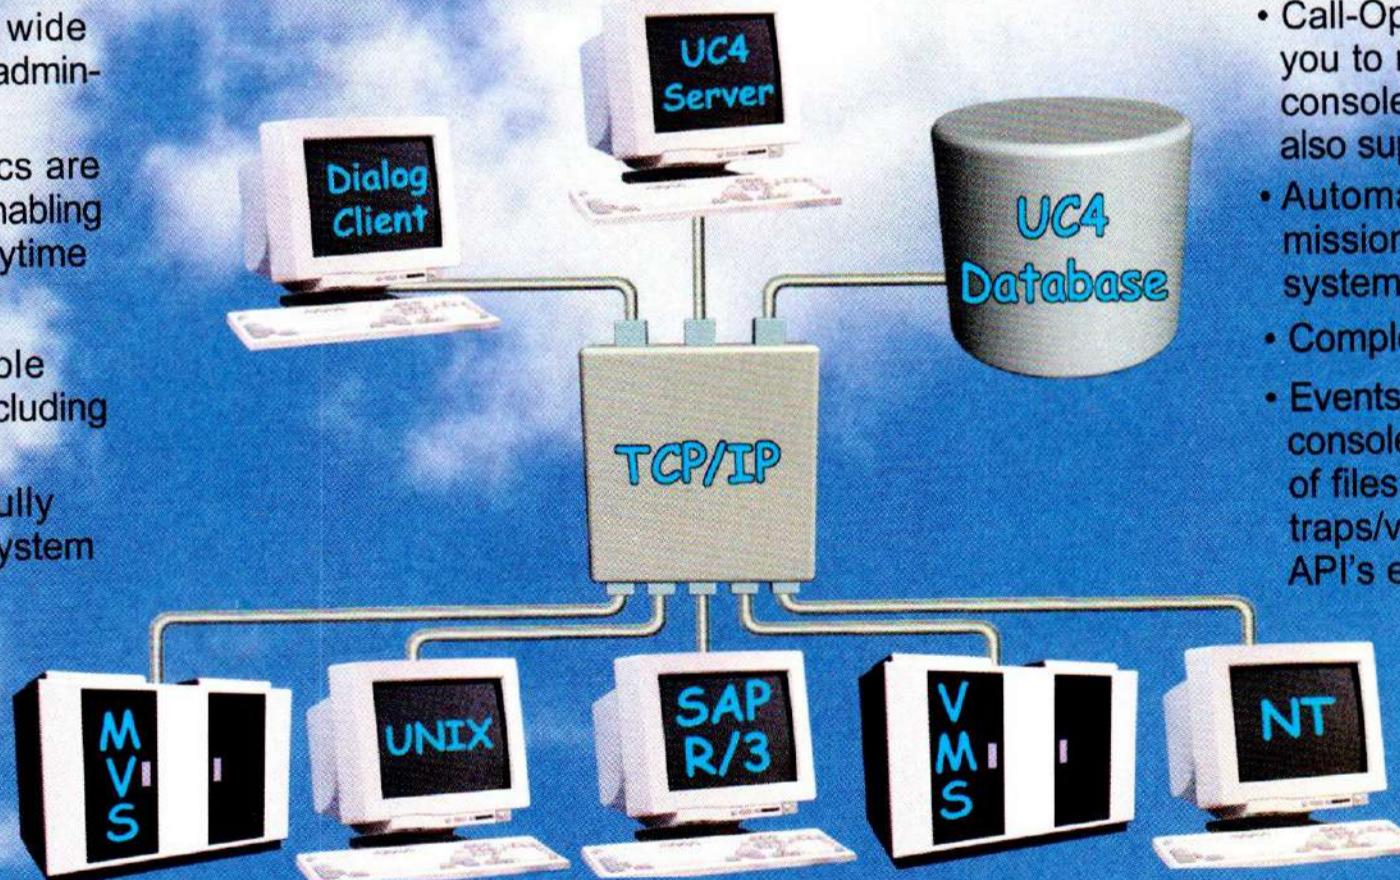
A Miller Freeman Publication

IBM'S RS/6000 HA-H70

"Enterprise Wide Job Scheduling"



- Single point enterprise wide control, supervision and administration of jobs
 - All activities and statistics are recorded in a database enabling you to review data at anytime
 - SAP Certified
 - Automated and traceable execution of R/3 jobs, including batch input sessions
 - SAP R/3 jobs can be fully integrated across any system or platform



Features and Benefits:

- Single Point of Control for all platforms
 - Supports international standards such as : ODBC, SNMP, TCP/IP, and SAP RFC
 - Graphically windows compliant eliminating extensive training
 - Global execution of jobs from a centralized location
 - Secured file-transfer between system platforms with code translation
 - Users can execute tasks without extensive knowledge of the operating system(s)

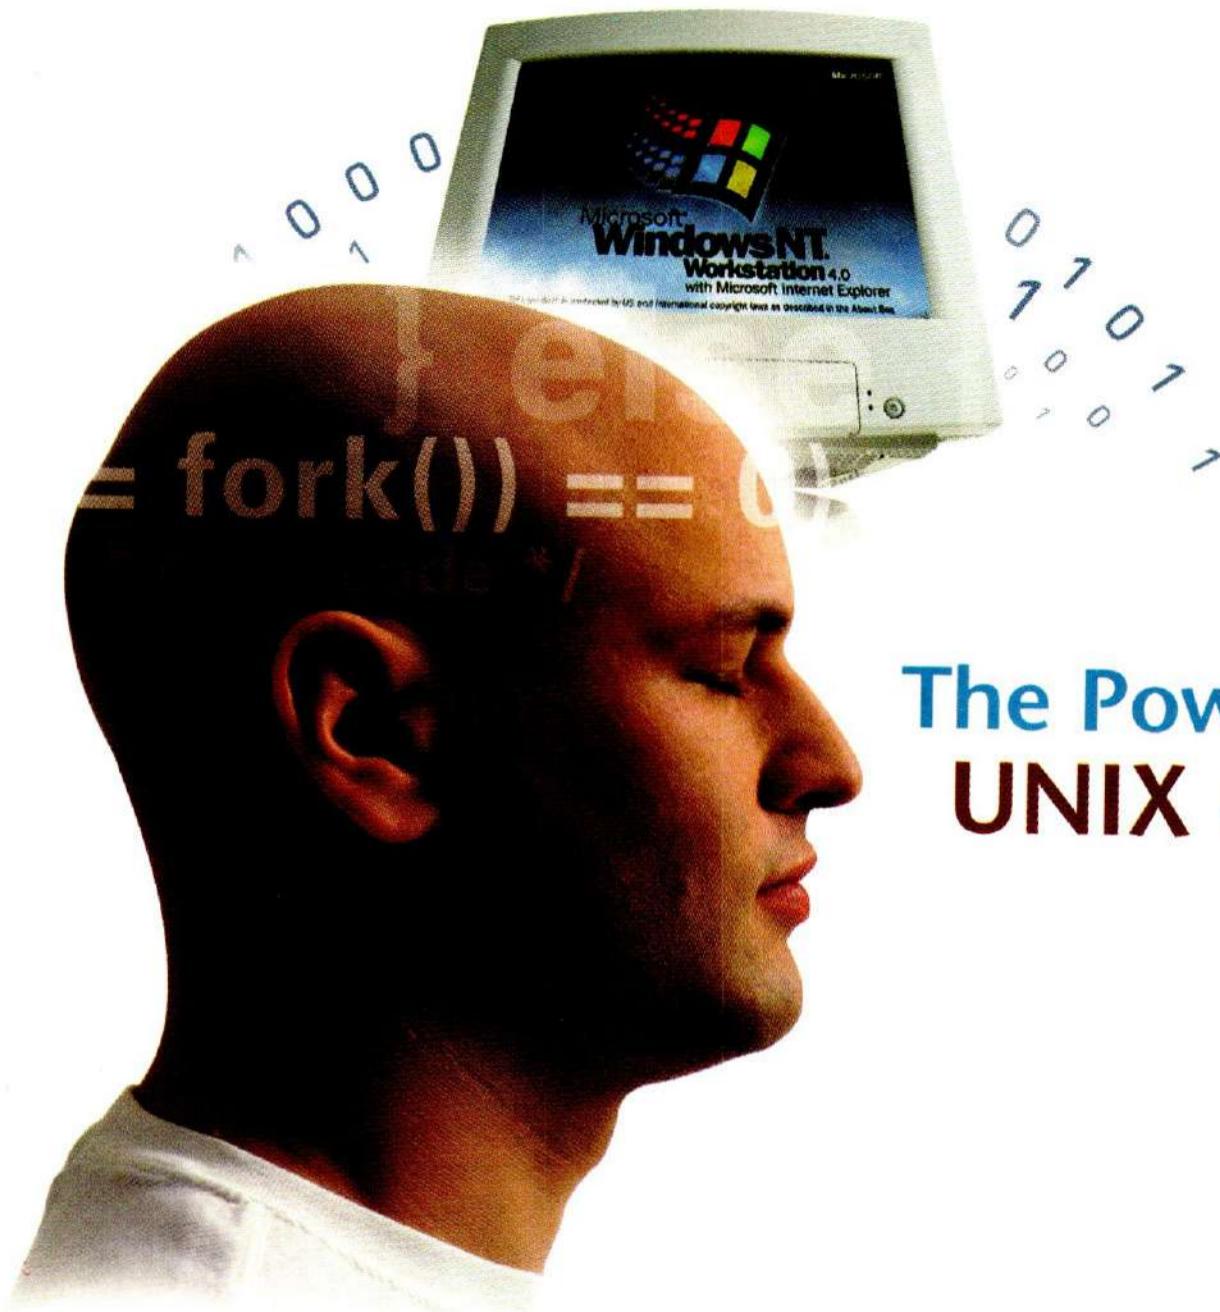
SEA™

SOFTWARE ENGINEERING OF AMERICA®

1230 Hempstead Turnpike, Franklin Square, New York 11010
(516) 328-7000 Fax: (516) 354-4015 Internet: www.seasoft.com

Select No. 190 at www.performancecomputing.com/fyi/

Call toll free: 1-800-272-7322



The Power of UNIX on WINDOWS!

NOW ALL YOUR UNIX KNOWLEDGE IS JUST AS USEFUL ON WINDOWS
...AND SO IS YOUR UNIX CODE!

Are you a UNIX® pro working on Windows®?
Developing? Deploying? Administering?

NOW ON WINDOWS YOU CAN:

- Use a UNIX command-line environment: shells, commands, and utilities
- Run UNIX scripts
- Port business-critical UNIX programs
- Do cross-platform development—write once & deploy on both Windows and UNIX
- Web-enable existing applications
- Interoperate with UNIX using X, NFS and telnet
- Integrate mission-critical and workgroup applications onto a common desktop

Choose MKS Toolkit®, Toolkit Select, or NuTCRACKER® Professional. All these products have been proven to cut learning, development and support time—70%, 80%, even 90%. And there are thousands of customers and case studies to prove it.

TO ORDER YOUR 30-DAY RISK-FREE TRIAL OFFER,
CALL 800-637-8034 OR 703-803-3343 OR
VISIT US ON THE WEB AT WWW.MKS.COM/INTEROP

ORDER YOUR FREE PORTING GUIDE TODAY!



MKS, MKS logo, MKS Toolkit, NuTCRACKER and NuTCRACKER logo are registered trademarks of Mortice Kern Systems Inc. UNIX is a registered trademark in the U.S. and other countries, licensed exclusively through X/Open Company Limited. Windows, Win32, and Windows NT are either registered trademarks or trademarks of Microsoft Corporation in the U.S. and/or other countries.

VERITAS is a registered trademark of VERITAS Software Corporation in the US and other countries. The VERITAS logo, BUSINESS Without Interruption and VERITAS NetBackup are trademarks of VERITAS Software Corporation in the US and other countries. Other product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.



Ever wish multi-platform network backup were this easy?

Presenting NetBackup, the simplest way to backup a complex network. It's designed to work with any combination of platforms, including Windows NT, UNIX and NetWare, as well as major databases and applications like Oracle, Informix, Microsoft SQL Server, Sybase, Lotus Notes, PeopleSoft, and SAP R/3.

NetBackup also gives you centralized control, and with Global Data Manager (GDM) you can centrally

manage backup and recovery operations anywhere on the planet. With surprising terabyte-per-hour backup speeds, and highly parallel, multiplexed backup and recovery you'll get the performance you need now, and in the future.

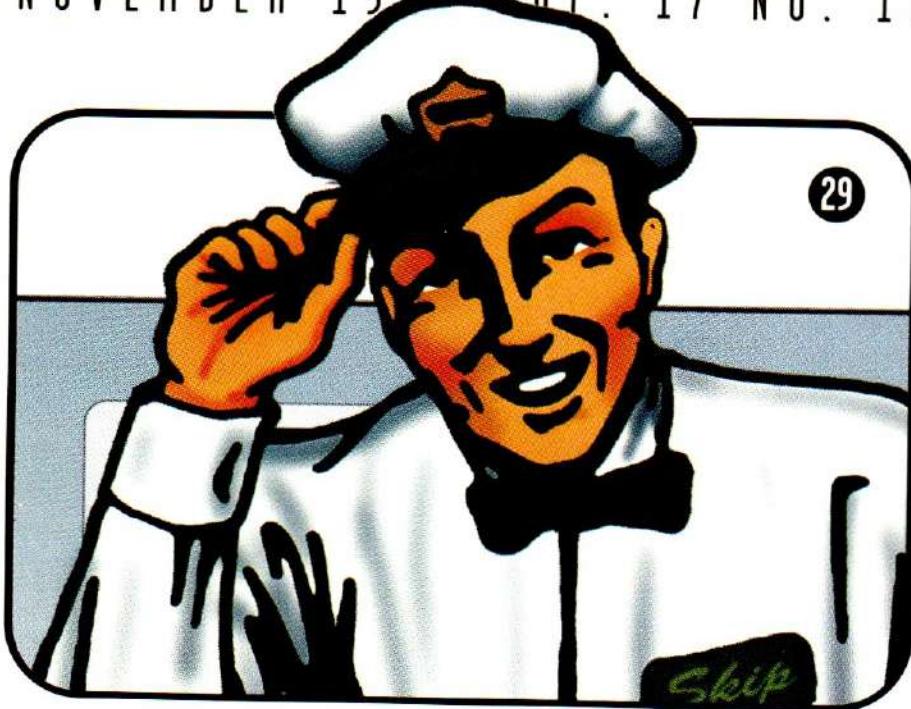
So call 1-800-729-7894, ext. 83517, or surf the web at www.veritas.com today. Because it's hard to imagine an easier way to backup a complex network.

BUSINESS WITHOUT INTERRUPTION™

VERITAS

PERFORMANCE COMPUTING

NOVEMBER 1999 VOL. 17 NO. 12



FEATURES

SPINNING AN OUTSOURCED WEB / 29

ISPs' systems-management procedures are an example of how to structure in-house methods, and what to expect if you outsource.

Joe Tier

REVIEWS

IBM'S RS/6000 HA-H70 CLUSTER SERVER / 39

You may not want to wear it on your wrist, but this 850-pound cluster keeps on ticking.

Ralph Barker

RED HAT LINUX 6.0 / 47

New features add business appeal, but enough for corporate clients?

Emmett Dulaney

ACCRUE INSIGHT 3.0 / 51

Accrue Insight 3.0 provides flexibility in measuring Web site traffic and analyzing visit patterns.

Ben Thompson

Cover photography by Carter Dow

PERFORMANCE COMPUTING (ISSN-0742-3136) is published monthly, with an extra issue in June, by Miller Freeman Inc., 600 Harrison St., San Francisco, CA 94107. Entire contents copyright 1999. Periodical postage paid at San Francisco, CA, and at additional mailing offices. POSTMASTER: Send all address changes to PERFORMANCE COMPUTING, P.O. Box 1280, Skokie, IL 60076-8280. For customer service, telephone toll-free 888-847-6188, outside the U.S. 847-647-8603. Subscriptions to PERFORMANCE COMPUTING are available at the following annual rates (13 issues): US\$55 in the United States, US\$60 in Canada, US\$75 in all other countries/surface mail, US\$119 in all other countries/air mail.

COLUMNS

DEVIL'S ADVOCATE: Rambling And Scrolling / 13
Stan Kelly-Bootle

INTEGRATION: Be Bop / 17
Tom Yager

JAVA ADVISOR: Java Grows On (Inheritance) Trees / 23
David Detlefs

DAEMONS & DRAGONS: Internet Black Holes / 55
William LeFebvre

CROSS THOUGHTS: High-Definition X Color / 58
Eric Foster-Johnson

Q&A: File Smiles / 72
Matthew Huff



DEPARTMENTS

FULL THROTTLE: How Far Out-Sourced? / 5

CONTRIBUTORS / 7

LETTERS / 9

ADVERTISER INDEX / 65

SPOTLIGHT / 62

RECENT RELEASES / 67

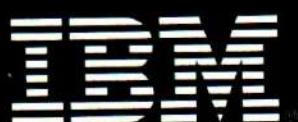


IBM, the e-business logo, DB2 and DB2 Universal Database are trademarks of International Business Machines Corporation in the United States and/or other countries. Windows and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both. UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group. Other company, product, and service names may be trademarks or service marks of others. ©1999 IBM Corp. All rights reserved.

DB2

little reasons to
check out DB2:
Palm OS and
Windows® CE.
Now the world's
most powerful
database runs on
the world's most
personal platforms.

DB2 Universal Database™ powers e-business solutions from the data center to the street corner, from palmtops to Windows NT® and UNIX®. Visit www.ibm.com/software/data.



411 Borel Ave., Suite 100
San Mateo, CA 94402
Tel.: 650-358-9500
Fax: 650-358-9739
<http://www.performance-computing.com>

EDITOR IN CHIEF
Ralph Barker
rbarker@mfi.com

MANAGING EDITOR
Jeff Swenerton
jswenerton@mfi.com

PRODUCTION EDITOR
Chris Munoz
cmunoz@mfi.com

WEB EDITOR
Claudia Willen
cwillen@mfi.com

ART DIRECTOR
Dwight Been

PHOTOGRAPHY
Carter Dow

DEPARTMENT ILLUSTRATIONS
Paul Watson and Audrey Welch

CONTRIBUTING EDITORS
Steven Baker, *Author*
David Detlefs, *Staff Engineer*,
Sun Microsystems Laboratories Inc.
Eric Foster-Johnson, *Author*
Glenn Groenewold, *Author/Lecturer*
Jon maddog Hall, *Executive Director*,
Linux International
Matthew Huff, *Consultant*,
Collective Technologies
Stan Kelly-Bootle, *Author*
William LeFebvre, *Consultant*,
Group Sys Consulting
Mary Petrosky, *Technology Analyst*
Kevin Reichard, *Author*
Randal L. Schwartz, *General Manager*,
Stonehenge Consulting Services
Tom Yager, *President*,
Applied Courseware Concepts

MILLER FREEMAN
A United News & Media publication

PUBLISHER
John R. Keough

VICE PRESIDENT/GROUP PUBLISHER
Katie A. Brennan

CHAIRMAN/MILLER FREEMAN INC.
Marshall W. Freeman

CEO/MILLER FREEMAN WORLDWIDE
Tony Tillin

PRESIDENT
Donald A. Pazour

EXECUTIVE VICE PRESIDENTS
Darrell Denny, Ed Pinedo (CFO),
Galen A. Poss, Regina Starr Ridley

SENIOR VICE PRESIDENTS
Andrew A. Mickus (Production),
John Pearson, Wini D. Ragus

Printed in the USA

<http://www.performance-computing.com>

FULL THROTTLE

HOW FAR OUT-SOURCED?

The use of the Web browser as a universal interface is gaining in popularity, particularly as a front-end to legacy applications that are hosted on large systems in the data center. While this trend usually makes the lives of users easier, the same cannot be said for system administrators and IT managers. Where many aspects of a business operation were once small, discrete operations that bore little connection with other functions in the organization, this is no longer the case. As workflows consisting of myriad small tasks run through a company's intranet and its Web servers, the task of managing this burden becomes more complex. Operational elements that once were trivial have become significantly more strategic when other workflows depend on them. Business inside the company is no longer done by memo and manual—now the company intranet has become the foundation of all critical paths, if not the path itself.

Three basic options now face the IT manager: maintain and manage all the systems in-house, outsource all or some portions of their operation to external service providers, or return to the era of tin cans and string. As the last option lacks both scalability and chicness, IT managers and systems administrators must choose between the former two.

As control of IT-related operations continues to centralize, internally managing all aspects of the company intranet certainly has appeal. Now that all those rogue PCs have been corralled by the new data center, it seems only natural for IT managers to clutch the gate keys tightly. All of the mission-critical operations are, after all, taking place on mainframes and

large servers that are managed by data-center personnel. Keeping intranet management in-house only extends the virtual corral out to the users' desktops.

Simple enough . . . if you have the staff. If you don't, there is little choice but to consider outsourcing at least a portion of the overall intranet-management task. How much to outsource will vary between companies. Some choose to outsource only certain Web-hosting functions, keeping more-critical services closer to home, under direct scrutiny. Other organizations may opt to outsource all but their desktop clients, relying on the ISP's IT staff to keep everything else running properly. Most savvy operations fall somewhere in between. The intranet-management objectives of these middle-ground companies are met by a mix of tightly worded service level agreements, and frequent monitoring of business processes.

To some, the concept of outsourced intranet management may seem strange in a country built on independence and self-reliance. But consider it a matter of hiring a data butler. The master of the manor can feel well served data-wise, even though the butler is out of town. The question is, how far out of town?

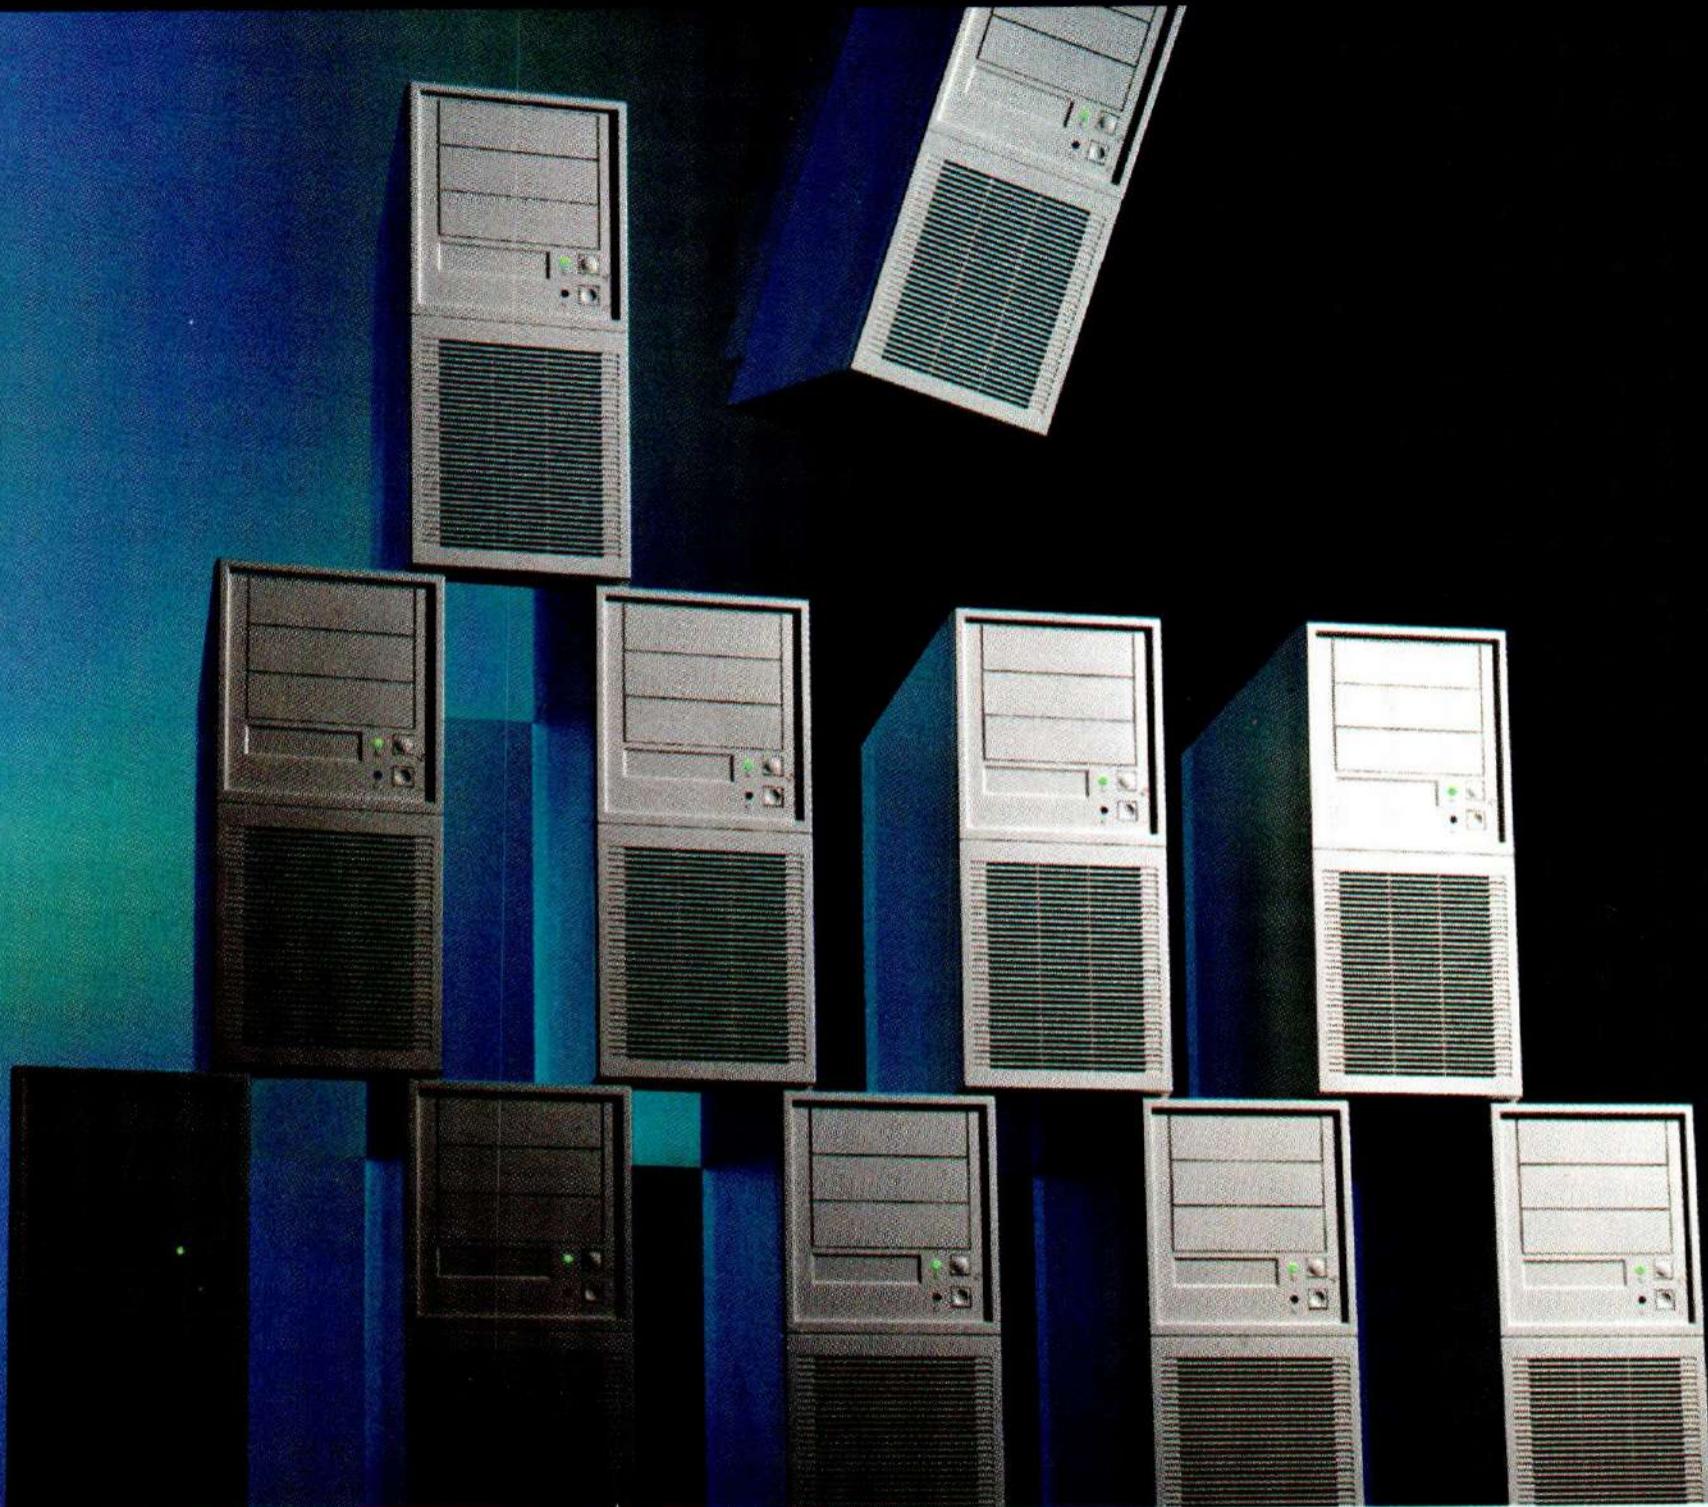


Ralph Barker, Editor in Chief

Will that next server make your network better?

Or just bigger?

© 1999 Intel Corporation. Intel and Pentium are registered trademarks, and Pentium III Xeon is a trademark of Intel Corporation.



The Intel® Pentium® III Xeon™ processor.

Don't just add a server—add the most manageable, scalable and powerful Intel processor-based server. When you're preparing your network for the challenge of e-business, there are no quick fixes. You need a long-term solution. One that provides more power without creating more complexity—the Intel® Pentium® III Xeon™ processor-based server. Performance and scalability combine to help simplify your e-business solution as you grow it. Built-in performance headroom means that when demand inevitably goes up, your system won't go down. Need extra capacity? Simple. You can scale from 2-way to 4-way to 8-way as your e-business needs dictate. And as you're adding capacity, you're adding simplicity, thanks to the processor's enhanced manageability. New challenges call for new solutions. And bigger isn't always better. Get your network locked, loaded and optimized for e-business—with the Intel Pentium III Xeon processor-based server.

[→ intel.com/IT](http://intel.com/IT)

intel.

CONTRIBUTORS

Ralph Barker ("IBM's RS/6000 HA-H70 Cluster Server," p. 39) is editor in chief of *Performance Computing*.

David Detlefs (Java Advisor, p. 23) is a staff engineer at Sun Microsystems Laboratories working on advanced JVM implementation techniques.

Emmett Dulaney ("Red Hat Linux 6.0," p. 47) is the co-author, with Vijay Sankar, of *Integrating UNIX and NT Technology: The Definitive Guide*, published by 29th Street Press. He is a consultant for DS Technical Solutions (<http://www.ds-technical.com>) and one of the original authors of *Inside UNIX* and several related titles. He can be reached at edulaney@quest.net.

Eric Foster-Johnson (Cross Thoughts, p. 58) is the author of *UNIX Programming Tools* and the primary author of *Power Programming Motif* Second Edition (with Kevin Reichard), published by IDG Books Worldwide. You can contact him at erc@pconline.com.

Matthew Huff (Q&A, p. 72) is director of operations for OTA Limited Partnership, a private hedge firm in New York, NY. He can be reached at mhuff@ox.com.



Stan Kelly-Bootle's
(*Devil's Advocate*, p. 13)
latest book is
Unix Complete, from
Sybex Inc.

Reach him at skb@crl.com or <http://www.crl.com/~skb/>.



William LeFebvre
(*Daemons & Dragons*, p. 55) is the editor for the SAGE series *Short Topics in System Administration*. He operates Group Sys Consulting (Alpharetta, GA) and can be reached at wnl@groupsys.com or <http://www.groupsys.com>.

Ben Thompson ("Accrue Insight 3.0," p. 51) is a Web systems analyst at Miller Freeman Inc., which operates a Web site farm for its own group of over 180 Web properties.



Joe Tier ("Spinning An Outsourced Web," p. 29) began his professional career at Allied Signal. Before entering the IS industry, Joe developed software for and tested air traffic control and electronic warfare systems. In the early 1990s, Joe worked on the development of the Data Warehouse at the Johns Hopkins Applied Physics Laboratory and eventually became a Microsoft Certified Systems Engineer. Joe is the technical director of operations at Digex Inc.

Tom Yager (Integration, p. 17) works at HealthWeb Systems Ltd. and is a freelance writer with a private research lab. E-mail him at tyager@maxx.net.

Highlights from

performance-computing.com

www.performance-computing.com



▲ **News accrues.** Our Recent Releases column is just too small to keep you abreast of all the latest high-end, data-center product announcements. So turn to our News section on [performance-computing.com](http://www.performance-computing.com) as your source for additional timely product information and major industry announce-

ments. The News column is updated regularly and indexed by date and product category. No matter what you're looking for, News is the place to look.



▲ **Linfo.** Check out the latest scoop on Linux and open-source variants of UNIX at [Linux-IT.com](http://www.linux-it.com). You will find articles such as "Embed Together: The Case For BSD In Network Appli-

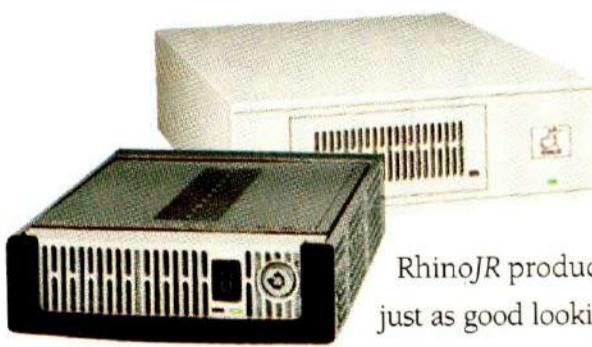
ances," by Kevin Rose and Charles Davidson, archives of Jon maddog Hall's "Penguin's Brew," and Linux-specific product information.

U100X

▲ **UNIX 100.** We're still counting down the Millennium, looking for the top 100 UNIX moments of the century. Go to <http://www.performance-computing.com> and click on the above icon to submit your pick for the best (or worst) of UNIX times.

**INTRODUCING RHINOJR.
A NEW ADDITION
TO THE
KINGSTON STORAGE FAMILY.**

www.kingston.com/storage



SOME FEATURES ARE SIMPLY INHERENT, as is the case with the new Kingston® Rhino®JR fixed and removable storage enclosures. Just like Kingston's rugged "Rhino" product line,

RhinoJR products are platform independent and constructed of sturdy metal. Nearly as tough, and of course, just as good looking, RhinoJR enclosures support fixed or removable implementations for today's fastest drives.

RHINOJR REMOVABLE ENCLOSURE

- Supports AT/IDE or SCSI Wide and Wide Ultra Devices
- (2) Auto-Sensing Cooling Fans with Audible Alarm
- 7-Segment SCSI ID and Drive Activity Display
- Power-On/Fan Failure LED
- Handle Assisted Carrier Eject and Key Lock

Call Kingston today at (800) 435-3973 to find out how you can adopt your very own RhinoJR.

RHINOJR FIXED ONE BAY ENCLOSURE

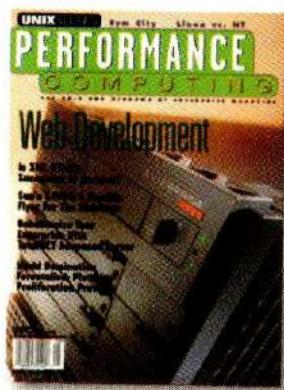
- Supports 3.5" SCSI 1, 2, 3, Ultra and Ultra2(LVD) Devices
- Auto-Ranging 40 Watt Power Supply
- Highly-Rated, Low-Noise 22.1 CFM Cooling Fan
- Drive Activity and Power-On LEDs
- Includes All Drive Mounting Hardware

Kingston
TECHNOLOGY
STORAGE PRODUCTS DIVISION



Kingston Technology Company, 17600 Newhope Street, Fountain Valley, CA 92708, USA (714) 438-1850, Fax (714) 438-1847. © 1999 Kingston Technology Company. All rights reserved. All trademarks and registered trademarks are the property of their respective owners.

Select No. 18 at www.performancecomputing.com/fyi/



Outsourcing to a non-American company or hiring a newly arrived programmer to the United States is a cheap way to get good labor.

—Chester Martel

H1B Re-visa-ted

I am amazed that you provided a full page to Daniel Brockman's diatribe [Letters, August 1999, p. 9] of how immigration laws "institutionalize discrimination" and "deny the human rights of liberty." Mr. Brockman represents himself as "President [of] Pacific Series Company" but fails to mention that he is the only employee of that company. Visit his "corporate Web site" at <http://home.earthlink.net/~brockman/>.

How is a nation putting limits on immigration any more of a human-rights violation than a corporation restricting the number of employees it hires? Are Sun and Microsoft guilty of human-rights violations for not hiring every qualified applicant?

According to the Department of Commerce, the United States will require at least 1.3 million new information-technology workers over the next decade (<http://www.ta.doc.gov/prel/pr06221998.htm>). Raising the H1B visa limit to 200,000 per year, as proposed, would add two million foreign workers. Clearly, Americans would be crowded out of the job market. Current laws allow corporations to fill positions by dialing "1-800-INDIA" without first recruiting U.S. workers. They can legally lay off U.S. workers while hiring and retaining "temporary" foreign employees in the same positions.

Brockman accuses American programmers of hypocrisy for condoning "discrimination based on geographic

origin." What a hoot. Perhaps Brockman can provide us with the list of U.S. corporations that pay the same wage to their programmers, regardless of whether they reside in San Jose or Bangalore.

KIM BERRY
CITRUS HEIGHTS, CA

I basically agree with Daniel Brockman's letter with one exception. He stated, "If software made by non-Americans reduces the cost of software, then the things produced by software will have lower cost." Balderdash, the American entrepreneur wants low cost and high profit. Outsourcing to a non-American company or hiring a newly arrived programmer to the United States is a cheap way to get good labor. Unlike Americans, they are willing to work long hours for peanuts, while their bosses become millionaires. Some American-based companies form divisions in foreign countries so they don't have to pay to bring the resources to the United States. The code is developed for nothing and then the software vendor has the mitigating gall to charge users based on an accelerated fee structure depending on the size of their computers.

CHESTER MARTEL

Our county recently increased the quota of immigrants we now allow to enter and work here. But that new



law is the thing that is discriminatory! This law allows hundreds of thousands more people into our country than before, the only hitch is that they must be technological workers. Non-technological people from foreign countries are being treated like criminals; as if they are not worthy of entry into our country. The law also promotes the brain-drain of foreign countries and erodes those countries' economies.

The real reason Congress is allowing more immigrant programmers in is to drive down the wages of American programmers. I don't believe these efforts will affect my income at all. Experience and excellent communication skills are necessary to create useful programs and interfaces.

Net effect: 0; moot point. Some employers will pay less and get less. Better companies will pay more and get more, just like before.

CHRISTOPHER J. INGRAM
SENIOR PROGRAMMER
BRANDT INFO. SYSTEMS INC.
TALLAHASSEE, FL

Daniel Brockman makes some interesting points in his diatribe against immigration limits, but his conclusions would only be valid if the "Tri-lateral Commission" imposed a "New World Order" successfully.

Brockman is correct that companies should not be allowed to discriminate based on national origin—I believe we already have at least a few laws to that effect. But this should not be accomplished by other forms of discrimination (for example, against older, more experienced Americans who demand reasonable pay for their work).

A lack of foreign programmers has not kept the cost of software high, as Brockman implies. With the exception of a certain popular operating system, I can think of no individual program (or suite of programs) I use that has not come down in price over the past ten years.

There is a difference between the programmer who moves from New York to Santa Clara, and the one who comes from China: the cost to our government, a burden that must be borne by the entire citizenry. Odds are, the programmer from China will create an unemployed American, not another job.

As for restrictions between borders, states have no power to restrict interstate trade, but that hasn't stopped most of them from trying. Every nation restricts trade to some extent; the United States already has the most liberal import laws in the world. We allow almost anyone to compete here, whereas we face substantial restrictions elsewhere. This is not going to change anytime soon, and Brockman better get used to that.

I understand why companies want to hire foreign workers—more often than not, they have greater skills at an earlier age than Americans. Our society has become complacent, if not lazy, and our education system is a joke. Perhaps instead of attacking the immigration system, Brockman should take on our failing education system. That would be an argument I could support.

JEFFREY L. FISHBEIN
RIVERWEB INTERNET COMMERCE
SELINGROVE, PA

Mr. Brockman's response to the article [Offshore Programmers: The Wave Of The Future?, May 1999, p. 14] is an interesting obfuscation of the real issue. As sentimentally reward-

ing as it may seem, appealing to the patriotism of American programmers will not change the basic facts—the software labor market is a market and, as such, is responsive to market forces.

He mentions greed and coercion but incorrectly fails to associate these actions with the corporate mindset. It was the "genius" of corporate wisdom, in its never-ending chase of profit, that broke the social contract with American labor during the 1980s. Be careful what you wish for. You may just get it—a temporary work force.

No employee that lived and tried to work through corporate downsizings, "rightsizings," and other labor blood baths will ever again subordinate their future to the mercurial temperament of corporate America. When executive salaries go to six and seven figures, and demographics show an ever-tightening labor supply, a reasonable person will take whatever advantage they can of the situation. That includes lobbying elected officials for competitive advantage in the labor market. Surely corporations do the same?

If it is a stable and loyal pool of labor that Mr. Brockman desires, then try earning your programmers' trust and respect. If you would rather have disposable employees, then don't expect any American to sit idly by while you import bodies to replace them in a country they helped build.

CLIFF WHEATLEY
PHILADELPHIA, PA

In Daniel Brockman's letter he suggests the primary reason to allow foreign programmers access to our markets is that of plain, old-fashioned democracy. It is interesting to note that Mr. Brockman chose to defend those countries that are accused of discrimination of the worst form. How many American programmers are allowed this type of entry into India or China? What about businesses? What do these countries do to safeguard intellectual property rights, international copyrights, and patent infringements?

If foreign programmers come to funnel jobs away from qualified applicants merely because they themselves

cost less; if they are here only to capture the intangible intellectual property of their employer and take it with them when they leave, if they are here to take advantage of a free society, then let's limit our exposure to them.

I agree that arbitrarily restricted markets make no sense at all. What kind of sense does it make when one side is open and the other closed? Liberty should be a two-edged sword, cutting both sides equally.

STEPHEN J. GUTHRIE
REGIONAL SALES MANAGER
MANTISSA CORP.
BIRMINGHAM, AL

Correction

The cover of the September issue incorrectly included a pointer to the "Gnome Gnuggets" article. The article is the Cross Thoughts column on p. 57 of the October issue.

My response to Cameron Laird's letter [Letters, September 1999, p. 9] attempted to combine too many elements of cache levels into too small a space, and completely confused the issue as a result. I thank those who sent e-mail advising of this error, and I apologize for the confusion.

The point that I was attempting to make was that while most OSs do a reasonable job of keeping active data in memory, circumstance and configuration can result in the data being moved to disk. In such cases where disk access is required to retrieve the data, the no-latency characteristic of SSD can be a significant win.

—Ralph Barker

Send your letters to:

- *Performance Computing*, 411 Borel Ave., Suite 100, San Mateo, CA 94402 or by e-mail to rbarker@mfi.com.
- All letters become the property of *Performance Computing*, are assumed intended for publication, and may be so used. Letters may be edited for clarity or space. Please include your contact information.

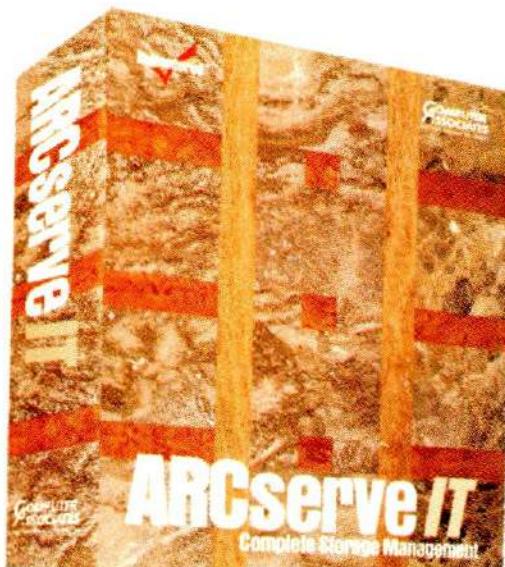
Everything You Need To Choose The Best Storage Solution.

Even though we're #1 in storage management and have received virtually every industry award in the category, we're not about to start resting on our laurels.

To the contrary, we're investing more money in R&D than ever before, and it shows.

With a broad range of exciting new features and benefits, ARCserve[®]/IT™ remains today's most advanced storage management solution, and now it's available in three editions to address the specific needs of your organization.

ARCserve/T **Workgroup Edition** protects your Windows NT LAN. With comprehensive backup for both server and attached clients, online backup of business applications and fast remote disaster recovery, this solution has everything a small environment needs.



Features	Computer Associates ARCserve IT v6.61*	Legato NetWorker v5.1	Seagate Backup Exec v7.2
Built-In Unicenter TNG [®] Framework [™] Management	✓		
Java-Based Remote Management	✓	✓	
Free Small Tape and Optical Library Support	✓		
Flexible Directory Level Open File Backup	✓		
Unattended Remote Disaster Recovery	✓		
Dynamic Tape Drive and Library Re-Configuration	✓		
High-Speed Image Backup	✓		
Free RAID Fault Tolerance	✓		
Shared Tape Library Support (ACSLS)	✓	✓	
Fibre Channel SAN Solution for Windows NT and NetWare	✓		✓
Unicenter TNG [®] Integration	✓	✓	
AS/400 Client Agent	✓	✓	
Complete MVS, OS/390 Storage Solution	✓		
Messaging System Protection			
▪ Microsoft Exchange	✓		
▪ Lotus Notes	✓		
▪ GroupWise	✓		✓

ARCserve/T **Advanced Edition** is a scalable solution that offers centralized management for medium to large environments that require multiple platforms and advanced technology support.

ARCserve/T **Enterprise Edition** provides comprehensive support for large-scale, more complex and demanding enterprise environments. With cross-platform administration, exceptional performance, and unmatched reliability, ARCserve/T



Enterprise Edition protects and manages your data from desktop to data center.

But don't let all of the power and advanced functionality fool you. ARCserve/T is still as friendly and easy to use as ever. Just what you would expect from the industry leader in systems, network and storage management.

To receive your FREE Trial Copy, visit www.cai.com/ads/arcserveit or call 1-877-2 GO FOR IT. Also available through CA's Open License Program™.

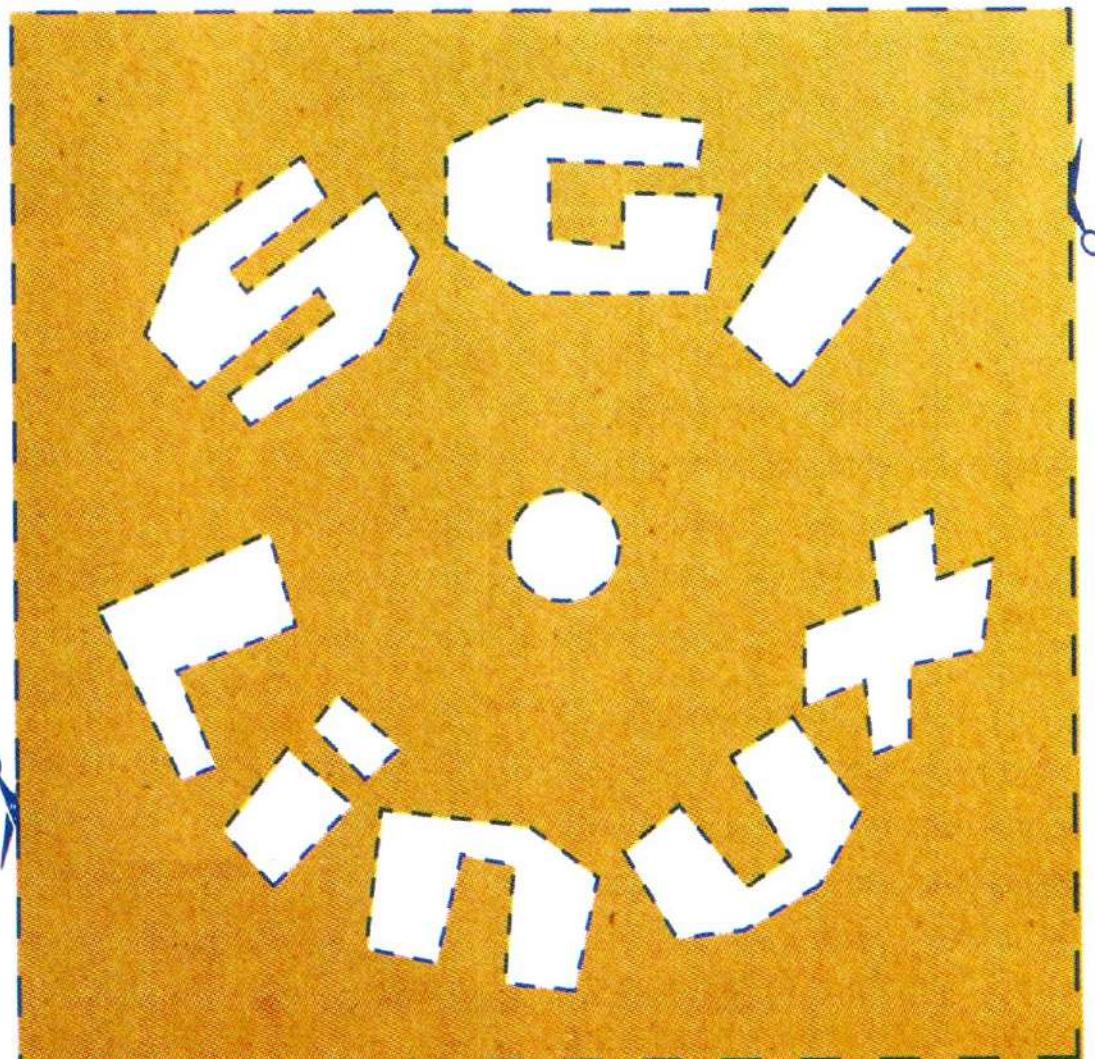
**COMPUTER[®]
ASSOCIATES**
Software superior by design.

ARCserve[®] /IT

© 1998, 1999 Computer Associates International, Inc., Islandia, NY 11788-7000. All product names referenced herein are trademarks of their respective companies. Information based on publicly available information as of 3/12/99. *Features vary by platform and edition, and may require additional options.

Rise Up.

- ① Buy an SGI™ 1400L server running SGI Linux.™
- ② Cut out stencil.
- ③ Spray paint everywhere.
- ④ Run fast.



www.sgi.com/go/linux 1•888•888•2772

The SGI 1400L Server.
Powered by SGI Linux.



The Pentium® III Xeon™ processor is specifically designed to enable server reliability.

©1999 Silicon Graphics, Inc. All rights reserved. Silicon Graphics is a registered trademark, and SGI, the SGI logo, The solution is in sight are trademarks, of Silicon Graphics, Inc. Linux is a trademark of Linus Torvalds. Intel, the Intel Inside logo and Pentium are registered trademarks, and Pentium III Xeon is a trademark, of Intel Corporation.

sgi
The solution is in sight.

Select No. 10 at www.performancecomputing.com/fyi/

Stan Kelly-Boote

RAMBLING AND SCROLLING

Every little once in a while one or more of my three original Muses stops making house calls. On a bad day I can spend hours debating whether the phrase “one or more of” requires a singular or plural verb form.¹ What would Fowler, Safire, Strunk, and Strank say? A blessed

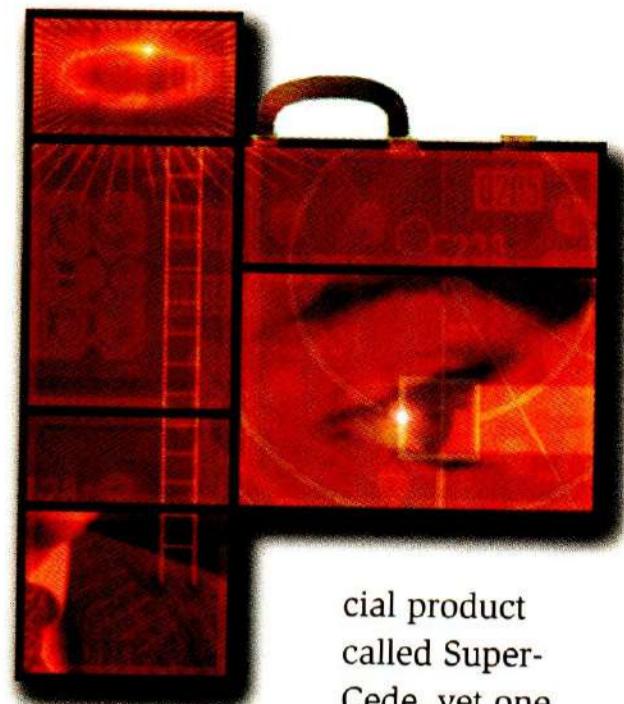
distraction ensues while I search the shelves and notice that my 40 Anchor Bible volumes are out of sequence. Lord, how could this be? And what is the *Bourne Conspiracy* novel doing in the UNIX section? Tom and Sam Becket[t] were both right: *Les choses sont contre nous!* They, them, some-ones, somethings are conspiring away out there, stemming my natural baignage [*You could have fooled me—Ed.*]. I’m no more paranoid than your average Secretary of State Nobel Peace Prize laureate. Yet, “I wonder who’s Kissinger now?”²

I sing for Aoede, meditate on Melete, and try to recall Mneme, but, damn their beautiful hides, they are busy elsewhere helping columnists other than *moi*. Why can’t a Muse be more like a *man*? The old cine-image of typewriter-blockage shows screwed-up balls of rejected pages piling up in the trash can, ashtrays overflowing, and the defenestration of empty bourbon bottles. Bretécher’s cartoons reveal the hung-up “authoress” washing dishes and rearranging the furniture, anything to avoid the daunting blank sheet. Of course, Joyce, Hemingway, Fitzgerald, Lowry, Breton, and that boozy lot, rarely faced the editorial challenge of “10,000 exciting, timely

words on Java servlets by 6:00 A.M. next Sunday.” And under the current text-generating regimen (the discontent of rich content), the screwed-up pages have physically disappeared only to survive fitfully and taunt future on-Web scholars as revs, baks, and diffs.

MAIL SAVE

My own writer-blockage is file-in-cheek: the Muses may desert me, but they continue to inspire my correspondents and fans. First, on the very reliability of Internettoyage, Bertrand (Eiffel Rules OK) Meyer reports disturbing errors in the published texts of several English and Russian classics, clearly due to scanning errors and lazy editing, and their subsequent persistification [horror-sic] on the Web. Meyer ran various AltaVista searches for those familiar “irresistible/able,” “computable/ible” and “supersede/cede” spelling challenges, to find an alarmingly high percentage of mistakes. For example, the “correct” *supersede* yielded 30,000 exemplars compared with 5,500 occurrences of the dubious *supercede*. He was careful to exclude references to the commer-



cial product called Super-Cede, yet one wonders if that corporate name-caller was aware of the Latin *cedere* and *sedere* roots? Note also that bland string searches would also register my columns (including this one), warning that *supercede* is etymologically suspect. Ensuing mail discussed the fact that if a spelling solecism exceeded, say, 20 percent of normal usage, both forms would find their way into a Webster or two. (Webster III, 1987, already says “supercede—see supersede”) Then there’s the Ridgeon’s horticultural firm selling SuperSeeds, but I digress.

There’s a welcome *ponticulus* here to the e-mail generated by my casual August-issue query: “Pete Becker and I are seeking a pithy Latin maxim that may predate the computer-linguistic ‘When the reason for a rule goes away the rule goes away.’” Rick Leir’s .sig offered the cynical “If a thing is done wrong often enough, it becomes right.” More germane was Rick Schnarr’s response from the lofty depths of the National Archives of Canada: “This appears to have been taken into French law, although I couldn’t find a citation pointing to Roman law. See <http://www.msh-paris.fr/red%26s/communic/adages.htm>, whence *Cessante causa cessat effectus, Plus de cause, plus d'effet, Argument logique: Cause génératrice, Effet—cessation*. Another version is *Cessant causa, tollitur effectus*. Later, Rick came up with *Cessante statu primitivo, cessat derivativus*—wonderfully relevant to current object-

oriented challenges: during garbage collection, we may inadvertently destroy "base" objects and leave dangling references thereto from "derived" objects. Or, maybe, vice-versa? No limit to that ancient classical wisdom. Rick quotes *Frustra [vana] est potentia quae nunquam venit in actum* as proof that the Romans were familiar with the problem of "provably unreachable" code.

BRIEF WIT

A welcome respite from those multi-page, bandwidth-depleting, multi-cc: "jokes": the simple one-line sign-offs, whether new or borrowed from faded T-shirts:

"Nuthin's certain, that's for sure!"

—Bob Williams, RLW Consulting

"My favorite 'e-' word these days is 'e-nough.'"

—Richard R. Kramer, Bel Air, MD

"No electrons were harmed in the making of this e-mail message."

—Jeff Colwell

And my current favorites from A. Non, the divinely dangerous: "It is now safe to switch off your computer," "If we shouldn't eat animals, why are they made of meat?" and "Bad file or command name. Nice typing though."

And the many unintended giggles: "Among the Dead Sea scrolls are many sectarian documents."

—discussion on KFAX radio, 1100 AM, January 14, 1998

"This saves everyone money by

elevating the traffic on the local Web site."

—a Cloudscape Inc. press release

"the first security middleware providing digital messages with the security associated with paper transactions."

—an E-Lock Technologies press release

One immediately thinks of the *Letters of the Zion Elders*, Hitler's diaries, and CyberGold's patent application.

"INCREASE YOUR WEB SITE'S SEARCH ENGINE RANKING. BE FOUND."

—From kgl204truck@skills.co.jp

Presumably, add Java-Sex-Teen-Gun-Blood-Java to your meta-thingsy?

REFERENCES

- Speaking of which, even the latest, greatest surviving MS-DOS (of uncertain version) running in or under (possibly over) the latest, greatest Windows OS (as in the clothing trade abbreviation for "oversized") still displays curious number-hedging messages such as "1 file(s) copied."
- USF Prof. Andrew Goodwin reminded me last Saturday during the 7:00 A.M. live-satellited Chelsea 4, Sunderland 0 soccer match that Kissinger stormed out of a 1973 BBC interview when the interviewer was gross enough to ask "Do you feel any shame in accepting the Nobel Peace prize . . . ?" History being what it is/was, this anecdote will no doubt be endlessly refried like Mexican beans: some will prefer to recall Kissinger overturning cameras and "purrin' a fluke's gob" on the entire BBC studio crew and audience. I should mention that many ex-pat-Brit soccer fans co-mingle at San Francisco's "Mad Dog In The Fog" pub on Haight at Steiner, former 1960s focus of the flowery peaceniks. While the U.K. yobboes are storming the terraces, we exchange polite polemics: "I say, the referee seems unduly tolerant to Beckham's consistently mistimed tackles."



Rave Computer Association, Inc.
Call Rave today at 1-800-966-7283 or
visit our web site at www.rave.net



Raves' total company commitment is to provide quality products and exceptional service.

Select No. 39 at www.performancecomputing.com/fyi/



Experience that “kid in a candy store” feeling.



When you see everything that the amazing Integrix line of servers has to offer, you may find yourself a little awe-struck, too. From web servers, to databases to firewalls, we have servers that can solve application problems no matter what kind of networking business you're in. And when you consider their compact design, UltraSPARC® processing power and density, our servers are the cost effective solution MIS managers are looking for. They're even built for NEBS compliance. And, as always, Integrix engineers will work directly with you to help achieve your network goals. Call us at **800-300-8288**. Or visit www.integrix.com to find out what a sweet deal Integrix servers can be.

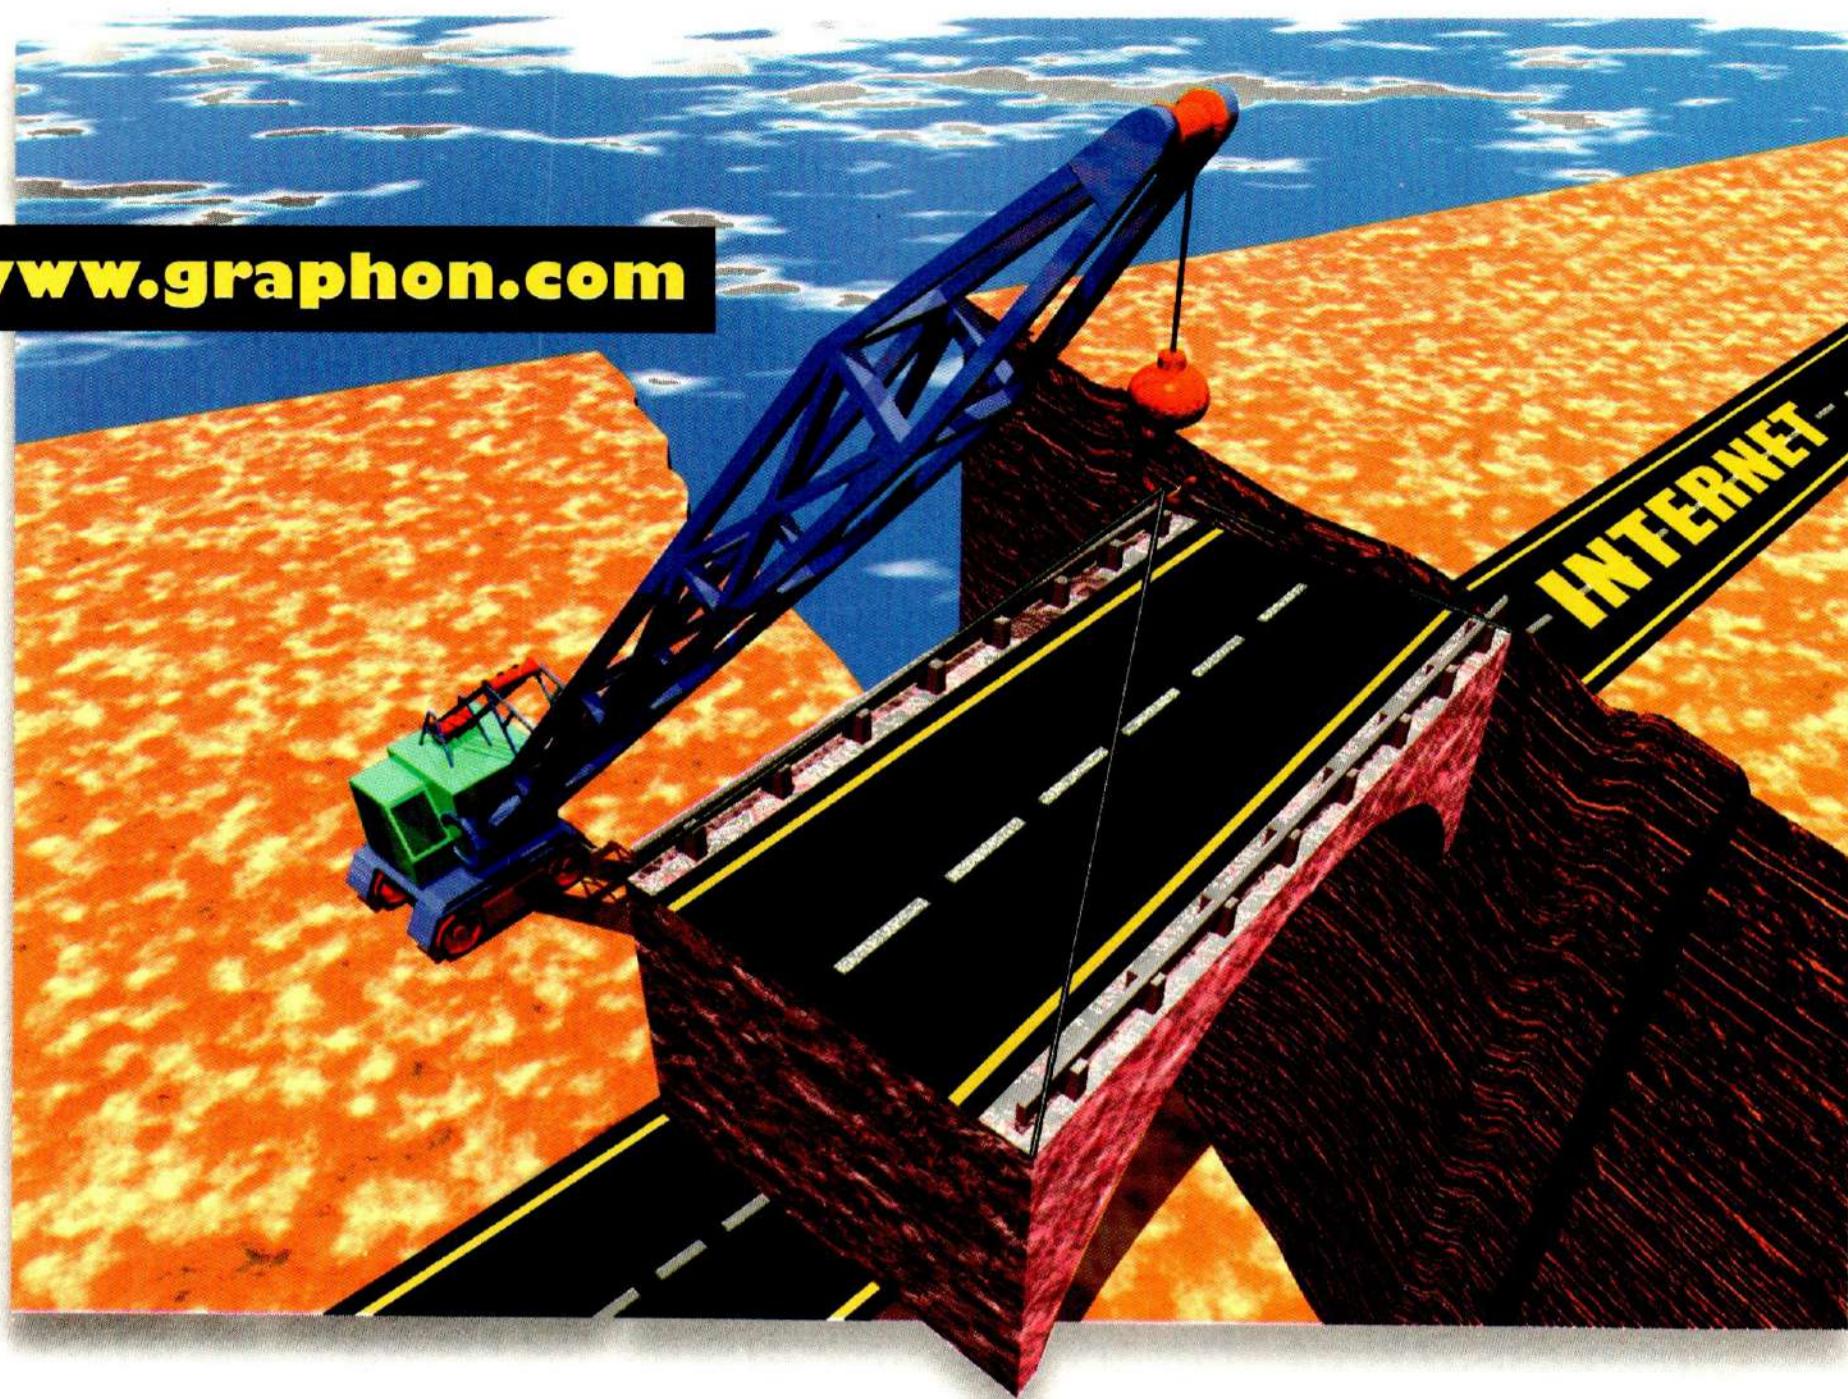
Select No. 140 at www.performancecomputing.com/fyi/



GO GLOBAL.

www.graphon.com

Oh yeah—one other thing. Power of Simple is a registered trademark of GraphOn. © 1999 GraphOn Corporation



**BRIDGE YOUR APPLICATIONS TO THE
INTERNET.
INSTANTLY.**

Any application. From anywhere. To anywhere. Over the Internet. Without software rewrites. No matter what hardware's on the desktop. With GraphOn. It's as simple as opening the package. Simple.

Simple for CIOs. Simple for ISVs. Simple for ASPs. Are you seeing a pattern here?

Did you notice that we said any application?

UNIX, Linux, NT. All simple. And completely interactive. In real time. And all applications work the way you expect them to.

To experience simple web-enabling, server-based computing, check out our website for a free demo. We'll show you the *Power of Simple™*.

Did you notice that we said free demo? Log on now.



Go global with web-enabling solutions that work.

Select No. 138 at www.performancecomputing.com/fyi/

Tom Yager

BE BOP

I've been covering the OS beat for about 15 years now, and in that time, I can recall seeing only three or four genuine innovations. It's been a while since something really spun my socks. BeOS, which I wrote off for years as a MacOS clone for PCs, recently crept back onto my

radar screen. This time I took a closer look, and was pleased and surprised by what I found. There are things I would change about it, but BeOS could be an alternative to Linux for those who are looking for the right mix of fast, cheap, easy, and reliable.

BeOS has been around long enough to be up to version 4.5 and it turns out to be pretty cool. I've seen many sound-bite summaries of BeOS that call it "Linux meets MacOS." Frankly, that spin doesn't grab my interest. BeOS had the misfortune of starting its climb at the same time Linux did, and the Linux people climbed a lot faster. The technical media lapped up the high-volume rhetoric from the Linux camp, which insists there is only one hip, alternative OS.

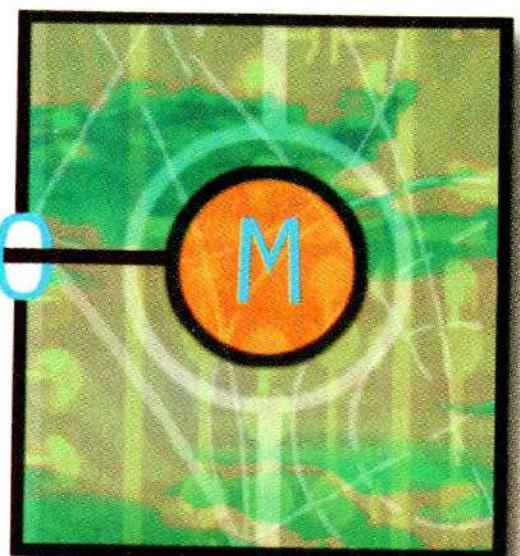
tem." It bridges several OS disciplines and has something to offer students, application developers, media (audio, video, graphics, and animation) producers, embedded-systems makers, and vertical-market software vendors. When you consider what each of these camps requires of an OS, you realize that no OS, not even the mighty Linux, has what it takes to serve all of their needs.

The world loves a good sound bite, so I'll spin a better one for BeOS: BeOS is a server-class desktop OS that isn't a server. It supports one user at a time (although other users may Telnet or FTP in). Its Web, FTP, and Telnet servers are anemic at best, which is appropriate

for a desktop. In a server vs. server drag race with Windows NT or Linux, BeOS would stall before the start. This fact, I believe, makes it hard for people to take BeOS seriously as a commercial contender.

True, we need Web, FTP, and file/print servers, but wouldn't a high-octane desktop be nice, too? While you consider your answer, mull over these stats: In my lab, BeOS 4.5 installed in seven minutes. It occupies about 150MB of disk space, including a full set of C/C++ development tools and UNIX commands. After the first boot, BeOS sprints from the boot menu to full GUI in 20 seconds, with no intervening text console, DOS, or UNIX underlayer. You boot straight to the BeOS graphical shell.

The BeOS microkernel is not, as many believe, derived from Carnegie-Mellon's Mach (as NeXTStep's is). Instead, BeOS's kernel was written from scratch, with inspiration drawn from the Xinu lightweight UNIX-like OS. The kernel supports up to eight CPUs and runs on Intel and PowerPC processors. If you wrote off BeOS as an alternative OS



OS PARADISE?

BeOS 4.5 contains many elements of an "ideal operating sys-

for Macintosh systems, you're in for a surprise. Apple has shut down its clone program and no longer releases the technical details of its systems. Because of Apple's return to proprietary platforms, BeOS cannot be adapted to G3 Power Macs. Be Inc. also no longer manufactures its BeBox system. Apple now has the only PowerPC system in the retail market. I expect the Intel platform will rapidly become BeOS's primary architecture. That's Apple's loss—BeOS does wonders for Intel-based desktop PCs.

The BeOS kernel has a scheduling latency of a mere 250ms. This makes it suitable for many real-time applications, and it improves responsiveness across the board. The rapid-cycling kernel, along with the applications and libraries that live above the kernel, are fully multithreaded. Each kernel thread will hop between CPUs in a multiprocessor system to balance load and keep time-sensitive tasks on track.

When you drive BeOS around, you can feel how well-tuned it is. It never seems to be working very hard. The user interface remains responsive until you push your system to its absolute limits.

READ THE FILE, WRITE THE FILE, BE THE FILE

When I approach an unfamiliar OS, I'm keen to know how it performs. Integer and floating-point benchmarks rarely turn up surprises unless the OS is an absolute pooch. These measure the speed of the hardware, not the OS. I find measuring an OS's disk I/O performance more telling. With a simple disk benchmark, I can open a window into the minds of the platform's designers. Windows 95/98, for instance, positively stink at disk I/O because performance wasn't a priority. Speed was sacrificed for broad compatibility and ease of use. Windows NT is about twice as fast as the lesser Windows, and Linux usu-

ally manages to move a bit faster than NT (depending on the disk controller).

Serendipity: I just booted my lab's Red Hat 6.0 Linux system to grab some fresh disk benchmark numbers. It took almost three minutes to fsck the server's 9GB Wide Ultra SCSI Seagate Barracuda drive. That is a fitting segue into a discussion of BeOS's file system. BeOS uses a proprietary, 64-bit, journaling file system. That it is proprietary shouldn't worry you too much, since BeOS will mount FAT and Mac HFS volumes for those times when compatibility matters. I mounted DOS FAT floppies, Zip, and Jaz disks without trouble. 64-bit storage addressing gives BeOS the ability to manage mind-boggling quantities of storage (millions of gigabytes). The journaling aspect of the BeOS file system provides instant recovery following a sudden failure. It takes no longer to boot BeOS after a dirty shutdown than after a clean one.

My impromptu BeOS disk I/O tests turned in the best numbers I've ever seen from my lab's desktop test machine. Using a Maxtor Ultra-ATA/33 drive, BeOS turned in a stable read speed of 13.7MB/sec. That is blazing, but not nearly as surprising as BeOS's write performance—which was exactly the same. I'm accustomed to seeing a 25–50 percent difference between read and write speeds. Writes are slower because the OS must locate free space and update file system structures as it writes the data. That is why you hear drives chatter even when you're writing a big sequential file to a clean drive. The heads keep seeking to the superblock (for UNIX derivatives) or the OS's version of the file allocation table. In contrast, the BeOS write was practically silent. The disk activity light was on, but except for occasional second-long bursts while the OS did something else, it made no sound. I should also mention that while the benchmark was running, CPU load rarely climbed above three percent. Other

applications were completely unaffected. My (rather unscientific) tests suggested that BeOS distributes its CPU and I/O loads well.

WRITING FOR BEOS

With the release of v. 4.5, Be Inc. turned away from the general desktop market (a wise decision; let Microsoft and Apple duke that out) and chose to beef up BeOS's digital-media capabilities. The visible result is an excellent set of software decoders for most popular formats of image, audio, and full-motion video files.

If you're a C++ developer, BeOS will feel like heaven. Be licensed a professional C/C++ integrated development environment (IDE) from Metrowerks Corp., makers of the CodeWarrior multiplatform C/C++ development system. You can buy CodeWarrior for Linux, Mac, and Windows and other platforms (even game consoles), but with BeOS, it's already in the box. CodeWarrior's editor performs syntax highlighting, brace, bracket, and parenthesis matching and lets you arbitrarily re-assign keys to functions. Double-click on a compile error and jump right to the offending code line in the editor. Run your application with debugging enabled, and a graphical debugger appears. If you need to look up something, the BeIDE, as it's called, will link you to the online documentation.

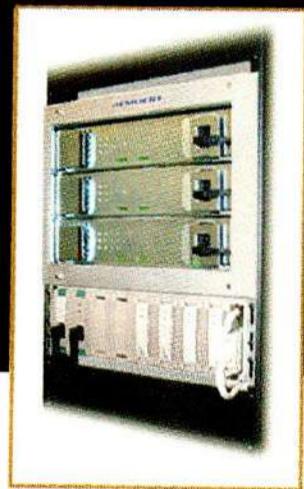
Be deserves praise for its choice, CodeWarrior is second only to Visual C++ as my preferred development environment, but Be's wooing of developers goes deeper than that. If you want to port code from UNIX or Linux, you'll be aided by the POSIX libraries and headers that ship with BeOS. The online documentation is vague on where the POSIX header files live, but once I found them I was able to port several UNIX test programs (including my disk I/O benchmark) with only a recompile.

When it absolutely positively has to stay up all the time.

Non-Stop Computing

360 MHz UltraSPARC

Firewall, DNS, Web...



Available in
Rack Mount

**Introducing the Resilient® Ultra fault tolerant server.
It delivers when any downtime or data loss is too much.**

We created the Resilient® Ultra fault tolerant server specifically for your mission critical network computing applications: firewalls, web, DNS, NIS+, license server, telecom, databases — any job where even a little downtime is too much.

Unlike software-based "high availability" schemes, our triple redundant architecture keeps applications running non-stop even if a logic board fails. Which means you won't experience unrecoverable data loss, service interruptions, high administration costs, reconfiguration delays — or constant worrying about your critical applications.

Resilience fault tolerant servers are based on industry-standard Sun UltraSPARC technology and the Solaris OS, so they integrate seamlessly with your existing network and applications. Their hot-pluggable modular construction offers exceptional ease of maintenance — you can even replace any failed component in under 60 seconds.

So now you can enjoy non-stop availability and the performance of an UltraSPARC server — at a surprisingly low cost. For an informative white paper with more information about the new Resilient Ultra, simply call us at (800) 977-9008 or visit www.resilience.com.

RESILIENCE
Fault-Tolerant Network Computing™

Select No. 178 at www.performancecomputing.com/fyi/

1755 Embarcadero Rd., Suite 120 · Palo Alto, CA 94303 · (800) 977-9008 · (650) 843-3600 VOICE · (650) 843-3609 FAX · email: sales@resilience.com

©1998 Resilience Corporation. Resilience, Resilient are registered trademarks of Resilience Corporation. UltraSPARC® is a registered trademark of SPARC International, Inc. Products bearing SPARC trademarks are based on an architecture developed by Sun Microsystems, Inc. All other trademarks are property of their respective holders.

In addition to the POSIX layer, BeOS offers its own interface to system services, but with a twist: most functions are exposed as C++ classes. To read and write files, you can use the POSIX-standard `open()`, `close()`, `read()`, and `write()`. If you opt for BeOS's object-oriented approach, you create a `BFile` object and manipulate the file's contents through method calls. When you destroy the object, the file is closed. Most system facilities, including BeOS's implementation of the OpenGL 3-D graphics library, are similarly object-wrapped. You're never prevented from using UNIX-style C APIs where they exist. After some time spent with the object-oriented approach, you'll be sold.

The BeOS user interface is what it is. It is neither revolutionary nor archaic. It mostly subscribes to the Mac's single-button mouse model, so Windows and X Window users' habit of right-clicking to pull up option menus will get them nowhere. The keyboard shortcuts are lifted straight from the Macintosh, with the PC's Alt key playing the role the Option key plays on a Mac.

BeOS's graphical interface handles text extremely well. It uses TrueType fonts, which are widely available. On a display with sufficient color depth, all text is automatically antialiased. Curved and angled lines are smoothed by interspersed pixels of intermediate (between the background and foreground colors) color. Linux and X Window users are tired of hearing this rant, but it's true: jagged-edged text looks cheap and unsophisticated. Ever since the Plus pack for Windows 95 added antialiasing to Windows, users have enjoyed cleaner, more readable text in their desktop applications, Web pages, and presentations. Every BeOS application inherits this feature, which should be standard on all graphical interfaces.

Console hounds won't be disappointed either. When you pop up a terminal window, the prompt you

see comes from the popular Bash (Bourne-again) shell. Bash is flanked by a full set of GNU commands and utilities. BeOS is not UNIX or a UNIX-derivative, but its POSIX support is quite complete and smoothly integrated with the rest of the system.

WHAT IS IT GOOD FOR?

BeOS's size, speed, and developer-friendly features make you want to write code for it. But is there any use for BeOS in commercial projects? I believe there is. There are plenty of shops that create custom software, and control the applications' deployment platforms. These are the developers who are using "specialty" OSs like OS/2 and QNX to develop embedded and vertical-market applications. Should they switch?

With a retail price of \$99, which Be has trimmed to \$69 for a limited time, BeOS is less costly than most specialty OSs. The development tools and object-oriented hooks make development fun and productive. Following deployment, BeOS takes care of itself; there is no administration to speak of. Its rapid boot time and journaling file system make BeOS a good fit for single-purpose applications like point-of-sale. While it is not marketed as a real-time system, BeOS's fast task switching and active CPU load balancing make it worth considering for timing-critical applications like industrial process control, high-speed data collection, and medical devices.

The strong digital-media playback support, high-quality text rendering, and simple user interface make BeOS a candidate for running self-contained kiosks. For that matter, I would favor using BeOS as a front-end for any commercial application hosted by a UNIX, Linux, or BSD server. Application vendors shouldn't make users deal with X Window interfaces when better alternatives exist.

BEWARTS

It's obvious I'm enamored with BeOS, but it isn't without shortcomings. The commands and utilities don't include familiar TCP/IP network tools like route and ifconfig, but you do get ping and netstat. Its Macintosh networking is seamless and thorough, but the early cut of Windows networking is unstable. A freeware NFS client is available. There is no NFS server for Intel-based systems and the SCSI tape backup is supplied by a freeware driver and utilities.

Commercial application support is scant, explaining Be's move away from the general-purpose desktop market. The digital-media niche is ripe for innovation, but Be will have to fight against many well-established players in a market that routinely makes mincemeat of newcomers. With its fast I/O, responsive user interface, and object-oriented media file support, it has what it takes to succeed. The risk is that digital-media hardware vendors, still in the midst of a costly migration from Macs to Windows NT PCs, may not want to risk resources on a lesser-known platform.

BeOS has its following, but its followers are nowhere near as rabid and religious as devotees of Linux and the Macintosh. Hardcore Mac users are insulted by the upstart Be, founded by a former member of Apple's inner circle. The Linux camp derides BeOS for not being open source. True, the OS's source code is not downloadable via FTP, and neither is the OS. The same is true of Oracle 8, DB2, StarOffice, WordPerfect, and all the other commercial applications Linux fans invoke as proof of Linux's success. There are users for whom open source matters, and others who just want the right tool for the job. Most users of networks on which I consult have no desire to crack open the source code for Word or Windows. Be shouldn't be lynched for protecting its trade secrets.

The range of supported hardware is narrow, which is another point in favor of BeOS's use in embedded and controlled vertical applications. This is not an OS you buy for a computer you already have. You might have lucked out and be using one of the five supported network cards, plus one of the nine supported graphics cards, but it's unlikely. It is more realistic to build a system based on BeOS's compatibility list. For embedded or pure front-end applications, get an ATX or micro-ATX Celeron motherboard, a CPU chip, a cabinet, a cheap ATI or S3 graphics card, an NE2000-based or Intel EtherExpress network card, 32MB of PC-100 RAM and a 2GB Ultra-ATA drive. A number of smaller companies sell preconfigured systems in this configuration or close to it.

Other reviewers and analysts have slammed BeOS for what it's not, invariably comparing it to Linux. In performing your own analysis, don't let yourself forget that BeOS is not a server OS. It is a fast and stable client with a great head for graphics and digital media. It speaks the language of C++ coders. There is one company-blessed book on BeOS, which is *The Be Book*, an HTML version of which ships with the system. BeOS development is a quick learn. You'll rapidly progress from relying on POSIX and the C runtime library to using BeOS's well-designed C++ classes.

One more dart some analysts toss at BeOS is its inability to run DOS, Macintosh, and Windows applications. It is the absence of support for non-BeOS applications that accounts for the OS's tiny footprint and ample speed. Microsoft and Intel should take note—Windows 2000 and the Pentium III would be better off if they gave up the ability to run DOS and 16-bit Windows programs. Let's close the book on being able to run Microsoft server OSs on 80486 CPUs. Those 30 users who want to run Word 2.0 on their PCs have my

blessing; I understand money is tight. Let them buy fire-sale 486s and Windows 3.1. The rest of us need to see Windows get smaller, and Intel make better use of its limited silicon real estate. Lack of PC and Mac application compatibility is not a liability for BeOS. It allowed Be's developers to focus on innovation.

I've appreciated other non-mainstream OSs, but this is the first time a desktop alternative held such appeal. Were I contracted to deliver a turn-key solution, I would make my choice for the server OS based on project requirements. Unless the customer had a compelling objection, I would put BeOS high on my list of candidates for front-end platforms.

- Build your own
- Shareware
- XRT Widgets for Motif

Professionals know there's only one answer to Motif interface development: XRT Widgets from KL Group.

XRT/table – Multi-purpose lists, tables and forms

- Data may be supplied programmatically, read from a file, entered by users, or supplied at run-time using a callback
- Supports table sizes of up to 2 billion rows by 2 billion columns
- Advanced Widget-in-Cell technology lets you efficiently manage widgets and pixmaps
- Includes label widget that supports rotated text and PostScript output
- Fast preview scrolling
- Royalty-free

XRT PDS also includes:

- XRT/graph – Wide variety of 2D charts and graphs
- XRT/3d – Sophisticated 3D charts and graphs
- XRT/gear – Essential add-on widgets and utilities
- XRT/field – Effortless data-entry and validation

XRT supports most major UNIX platforms, VMS, and Windows via Interix and NutCRACKER. Integrates with major GUI builders including UIM/X, Builder Xcessory, TeleUSE, View.h++, X-Designer and XRunner.

FREE XRT Widget Tips and Progress Bar at:

www.klgroup.com/widgets



KL Group Inc.
1-888-328-9591
Tel: 416-594-1026
Fax: 416-594-1919

KL Group Europe
Tel: +31 (0)20 679 95 03
Fax: +31 (0)20 470 03 26



XRT, XRT PDS, XRT/gear, XRT/graph, XRT/field, XRT/table, XRT/3d and XRT/graph are trademarks of KL Group Inc. All other products are trademarks or registered trademarks of their respective companies. PC000

Select No. 7 at www.performancecomputing.com/fyi/

SAVE DATA!

A sepia-toned photograph of a steam locomotive pulling a long train through a landscape. In the foreground, a person wearing a flight suit and helmet lies face down on the ground, appearing to be in distress or unconscious. A speech bubble above them contains the text "Help! My name is Data.".

Get on the right track with AMASS

Relax. There's no need to panic. Making the right choice is easy with AMASS® Online Software. AMASS is engineered so you can save data by extending storage capacity using an automated library. This process is transparent and appears as a single drive letter on the user's system. Because AMASS provides infinite storage, not just archiving, you can save it all.

Help Save Data!

Visit www.adic.com/savedata3 to learn more and register to win a flight bag to store your stuff.

adic

1 800 336 1233

David Detlefs

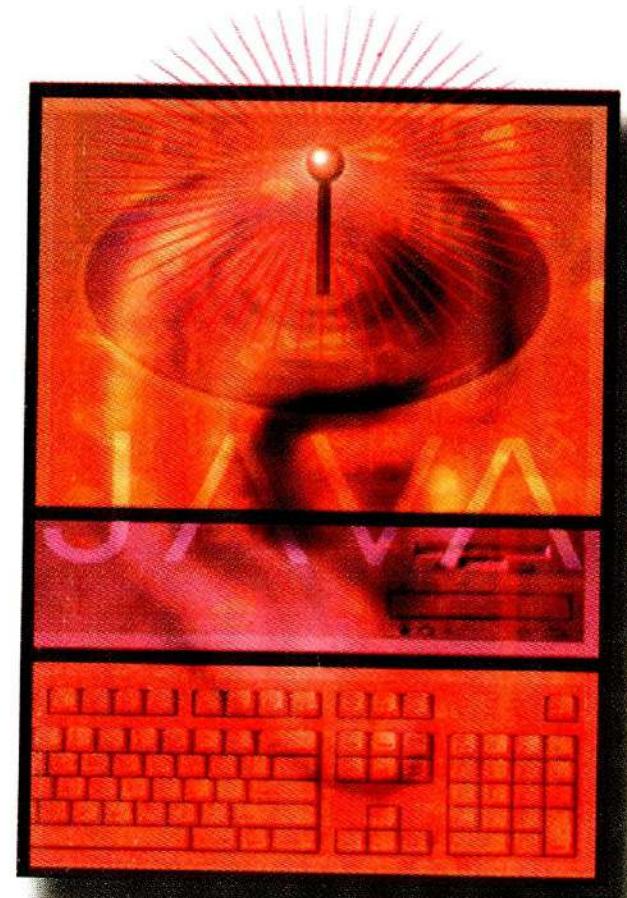
JAVA GROWS ON (INHERITANCE) TREES

Most textbooks on object-oriented programming tout the wonders of class hierarchies—collections of classes related by subtyping. Read enough and you may wonder how anything ever got done without them.

However, my experience is that if you're a typical programmer, you won't often have occasion to design such hierarchies. This is good, because it's a difficult thing to do well. But you will often have occasion to use them, so it's important to understand how they work. In this column, we'll examine a specific hierarchy: the classes of Java2's `java.io` package. These form a moderately complicated inheritance tree; by mapping out this particular design we'll learn something about such hierarchies in general.

With a couple of minor exceptions (the array-store check—yuck!), Java is statically strongly typed. This means that if a method expects an argument of type `A`, then the compiler will make you provide an `A` argument at every call. Instances of a class or interface `B` that extend a class or interface `A` are considered members of the type `A`, and instances of a class `D` that implement an interface `I` are members of the type defined by `I`. These are just syntactic rules: if I define a class `A`, with specifications on how its methods are to behave, you could certainly create a subclass `B` of `A` that overrides `A`'s methods so that they no longer obey those specifications, and there's little that the compiler can do to prevent it. Suffice it to say that this is a really BAD idea! If client code is written to expect an instance of type `A`, and relies on the methods of `A` behaving according to their specifications, then that client code will break when given a misbehaving `B`.

The foregoing is a pretty basic review, of course, but understanding it is nonetheless important to the design of class hierarchies. Classes or interfaces near the root of the hierarchy, which many classes extend or implement, must be carefully designed and specified: much client code will be written knowing only that objects are of these general classes, so the specifications of the methods of these classes must be sufficiently detailed to let the client code



do useful work. On the other hand, these specifications must not be overly detailed, or else they will restrict the set of possible subclasses.

An important question in Java is whether general types near the root of the hierarchy should be expressed as classes or interfaces. Many experts argue for the use of interfaces whenever possible, and I don't disagree. Interfaces simplify matters because they have method signatures and specifications, but no state. When a class, which may have state, is designed for use as a superclass, we always have to remember that such a class has two varieties of users: clients accessing instances of the class using its public view and subclass methods accessing the class's state (recall that fields may be declared as protected to make them accessible to subclasses, but not to arbitrary clients of the public view). For classes near the root of an inheritance hierarchy, it is as important to design this subclass view as it is to design the public view. This is a tricky business, which is why it is usually preferable to use interfaces, which have no such distinct subclass view. Sometimes, however, using classes with well-designed subclass views allows implementation commonalities useful to many or all subclasses to be factored out into the superclass—an irresistible proposition. We will see examples of this in the `java.io` hierarchy.

Actually, there are several different hierarchies within `java.io`. The classes `Reader` and `Writer` and their subclasses handle input from, and output to, streams of (Unicode) characters. The classes `InputStream` and `OutputStream` are similar, except that they read from and write to byte-based streams. There is a small interface hierarchy: `DataInput` and `DataOutput` are implemented by byte streams that provide operations to read and write the primitive types of the Java Virtual Machine (JVM), and the `ObjectInput` and `ObjectOutput` interfaces extend these (respectively) to reading and writing entire objects. (These last two stream types are used in object serialization, one of the foundations of the Remote Method Invocation system.)

Let's look at `Reader` and its subclasses in more detail. `Reader` is an abstract class, meaning that it has one or more abstract methods—methods that have no implementation in the class that defines them. Abstract classes may not be instantiated; only non-abstract subclasses that implement all the abstract methods are allowed to be allocated. `Reader` has only one field, named “lock”—the object that should be locked to synchronize access to this reader. Usually one would synchronize access to a data structure by using synchronized methods, which lock the object itself. And indeed, the constructor of no arguments initializes the lock field to the current object, “this.” The ability to specify another lock object is used by some `Reader` subtypes that delegate operations to another reader. For example, a `FilterReader` interposes some filtering operation on another reader; `FilterReader` specifies the other reader object as its lock object.

As you might expect, `Reader` provides a `read` method, in fact, several forms of `read`. The most general version attempts to read characters into a specified subrange of a character array, returning the number of characters actually read:

```
abstract public int read(char cbuf[], int off, int len)
    throws IOException;
```

This method is abstract: this is what subtypes must implement to become interesting Readers. The other two variations of the `read` method illustrate two useful idioms in class hierarchy design. The first is the method

```
public int read(char cbuf[]) throws IOException {
    return read(cbuf, 0, cbuf.length);
}
```

This simply defines a convenient shorthand for a common usage pattern of the more general `read` method. Note that the implementation can be provided in the supertype, even though the method called is abstract in that supertype. The implementation of this shorthand `read` can be inherited without modification by all subtypes (and causes no loss in efficiency, if your JVM's just-in-time compiler does a decent job of inlining).

The last version of `read` yields a single character (as an integer, so that -1 can be used to indicate end of stream).

This case is somewhat similar to the one above, in that `Reader` provides an implementation in terms of the more general abstract `read` method, but it is different in that the implementation is not especially efficient: a special-purpose single-character `read` implemented in a particular subclass can probably do much better. Thus, a subclass for which efficient single-character reading is not considered important inherits a version of the method that is correct but inefficient; subclasses for which such access must be efficient can override the method appropriately. It is generally useful to attempt to provide such “generic” implementations where possible in root classes of inheritance hierarchies, since it decreases the work necessary to get a working subclass.

The `skip` method, which reads and discards a specified number of characters, is another example where a general implementation can be provided in the superclass, since the method relies only on calls to the abstract `read` method that subclasses must implement. Here the general superclass implementation probably is sufficiently efficient for most subclasses. Here is the class hierarchy rooted at `Reader`:

```
Reader
  +-- BufferedReader
  +-- LineNumberReader
  +-- CharArrayReader
  +-- FilterReader
  +-- PushbackReader
  +-- InputStreamReader
  +-- FileReader
  +-- PipedReader
  +-- StringReader
```

We can divide the subclasses of `Reader` into two categories. The first category consists of `Reader` subclasses that correspond to different sources of character streams. `StringReader` and `CharArrayReader` allow the characters in strings and character arrays to be accessed using the general `Reader` functionality. `InputStreamReader` represents a reader whose source is an `InputStream`. Remember that readers yield Unicode characters and streams yield bytes; an `InputStreamReader` is a reader that translates the bytes in an input stream into characters, according to some encoding. There are many possible encodings; the constructor of the `InputStreamReader` specifies which is to be used. A `FileReader` is an `InputStreamReader` whose `InputStream` source is a `File`. Finally, a `PipedReader` is a reader whose source is a `PipedWriter`. Characters become available from the `PipedReader` when they are written to the `PipedWriter`.

The second category of `Reader` subtypes “wrap” extra functionality around existing readers. For example, a `BufferedReader` allows buffered access to some other reader. Buffering allows efficient character-at-a-time access to reader types for which such access might otherwise be prohibitively expensive, such as disk files. The `BufferedReader` maintains an array of characters that are read an array-full at a time from the source reader; read requests on the `BufferedReader` are



The holidays are coming...what to get those precious grandkids? Maybe some futuristic, high-tech gizmo. Or the latest craze. Then suddenly you flash on your favorite Lionels and you know the time is right

So you opened the carton where they've been for years and painstakingly cleaned every part.

You rewired, added new track and bought some cool accessories. And when at last you pushed the throttle, the locomotive started smoking...better than it used to.

Ah, the toys of childhood...lovingly restored to absolute perfection.

A gift they'll treasure.
A family tradition.

That's how we feel about computers.

New isn't always better.

Get on the fast track to performance and savings

Many of today's leading companies are using our Repowered computers every day in mission critical applications. Because they perform as good — or better — than new. Because they are backed by a totally risk-free guarantee. And because the savings are substantial...regardless of your current discount! So make tracks to our website and discover the savings for yourself.

Click on www.radiantresources.com for aggressive pricing, net specials - and the chance to configure and instantly price systems that will travel with you into the future.

Look who's Repowering with Radiant

If you want to know how reliable our computers are, just ask the companies using them in production environments — large telecommunications companies, internet service providers, brokerage firms, online stock trading and quotation services, web hosting and development companies... and, would you believe — a disaster recovery company!

So don't get railroaded into always thinking new. 'Cause like the workhorse performance of our 1937 Lionels, new isn't always better — especially when it comes with our better-

than-new, best-in-the-industry Radiant guarantee. And a high-speed ticket to spectacular savings!

Call Toll-Free—800-500-7675



Radiant
Resources, Inc.

**Repowered Computers
For The Next Generation**

1913 Atlantic Avenue, Suite R-3
Manasquan, NJ 08736
973-442-5555
Fax: 973-442-5550
www.radiantresources.com

Major Credit Cards Accepted.

satisfied from this buffer array until it must be refilled from the source. A `LineNumberReader` augments `BufferedReader` with one extra feature: it counts line breaks, allowing the current line number in the reader to be queried. `FilterReader` is an abstract class that represents any reader that uses another reader as its source, adding some extra functionality. `FilterReader` doesn't provide much functionality beyond exposing the existence of the source reader; on the other hand, it doesn't constrain its subtypes much, either. The only subtype of `FilterReader` provided in the `java.io` package is `PushBackReader`, which allows characters to be "unread," and pushed back onto the input stream. This functionality can sometimes be quite useful in language parsers, when some grammar requires look-ahead to disambiguate constructs. (Usually, however, the stream allowing the "pushback" functionality is a stream of lexical tokens, for example, keywords, constants, or identifiers, produced by a lower-level lexer.)

It could be argued that `BufferedReader` should have been a subclass of `FilterReader`. After all, it fits the general definition of `FilterReader` given above: it adds some extra functionality to a source reader, in this case, efficient fine-grained access to the characters. And I know people who will argue points like this at length, especially if the beer is cold and the

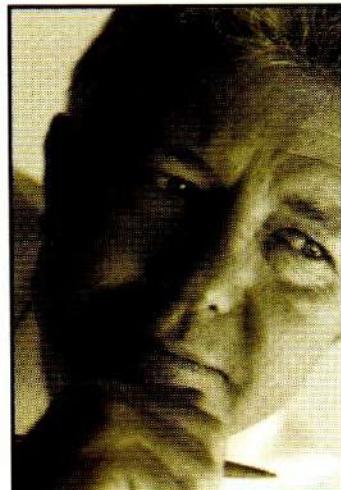
peanuts salty. To my mind, however, it's kind of a silly point to argue over. The right perspective from which to evaluate such arguments is "Does either alternative have an advantage over the other in efficiency, maintainability, or ease-of-extension?" I'm trying to emphasize relatively objective criteria over more subjective attributes such as "conceptual coherence." To evaluate maintainability or extensibility you have to imagine what bugs might be found, or extensions desired, and consider how much code you'd have to modify to fix the bug or implement the extension. If two designs aren't distinguishable in these objective measures, then my guess is that they're not all that conceptually different either.

In the particular case we're discussing, `FilterReader` adds so little functionality (exporting to subtypes only a single protected field identifying the underlying source reader) that duplicating that functionality, as `BufferedReader` does, is no great crime. If one were to create a reader subtype that did its own buffering, however, this would be felony programming: when someone else has done complicated programming like this, it is your duty to reuse that work. Unless you can really prove to yourself that it's not suitable for your purposes—then it's your duty to see if you can make it suitable. Reimplementing it yourself should always be a last resort.

The delegation idiom represented by `FilterReader` is an important one, and not just in the Reader hierarchy. In some ways, this is the Java language's best analog for multiple inheritance of implementations. GUI toolkits are often used in these examples; some languages support the construction of, say, `TitleBarredScrollBarbedWindow` by multiple inheritance from `TitleBar`, `ScrollBar`, and `Window`. A single class results, supporting each of the protocols defined by the superclasses, and inheriting implementations from these superclasses as well. (Sometimes some amount of coordination code must be provided in the inheriting class.) Java does not support inheritance from multiple classes, but does support classes that implement multiple interfaces. To get an effect similar to multiple inheritance, then, we'd have `TitleBar`, `ScrollBar`, and `Window` interfaces; the class `TitleBarredScrollBarbedWindow` would implement all three. But this class can choose at most one class as a superclass from which to inherit implementation. Let's say we choose `Window`. To get the effect of multiple inheritance of implementations, then, we might have classes `TitleBarImpl` and `ScrollBarImpl` that implement the corresponding interfaces. We can then program `TitleBarredScrollBarbedWindow` to have fields containing references to instances of these classes, and implement the methods of the `TitleBar` and `ScrollBar` interfaces by delegation to the corresponding methods of these instances.

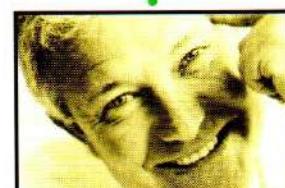
The rest of `java.io` has similar structure. `InputStreams` provide a hierarchy very similar to the reader hierarchy. (Interestingly, here `BufferedInputStream` is a subclass of `FilterInputStream`.) The output classes `Writer` and `OutputStream` are analogous to the input classes we have discussed.

I hope this has been a useful tour of `java.io`, and has illuminated some issues and principles in the design of class hierarchies.



- ? Got character-based UNIX apps?
- ? Sick of the Green Screen Effect?
- ? Tired of lost business, low productivity?
- ? Want GUI, multimedia and web connectivity today?

We've got you covered!



- A totally unique process
- Customizable, turn-key solutions
- Immediate results
- No Programming

Participate in the JSB Challenge!

You have nothing to lose!

Visit us today at www.jsb.com/challenge

Or call us now at **800.572.8649**

Free evaluation CD

JSB Corporation

108 Whispering Pines Drive, Suite 115, Scotts Valley, California 95066, USA.
Tel: 831-438-8300 Fax: 831-438-8360

Select No. 99 at www.performancecomputing.com/fyi/



—Burning Your Operating Systems at Both Ends?—

Let's Face it: integrating UNIX and Windows NT is an immense challenge. You may even be ignoring the migration problem, hoping it will go away. To prevent disaster, you need a resource to guide you through the solution-seeking process.

unixintegration.com is a Web site focused exclusively on helping you integrate and manage your UNIX and Windows NT systems. To help keep you on track as you merge these two complex environments, this site includes a monthly article that focuses on UNIX/NT integration, along with pertinent articles from **Performance Computing** and **Windows NT Systems**. Topics include sharing e-mail resources, scalability, the Internet, and database compatibility.

This site also features an Integration Q&A section that answers pressing technical questions, an Integration Library that guides you to additional resources, and a Product Gallery that gives you an instant connection to pertinent vendor information.

— Visit www.unixintegration.com today! —

:unixintegration.com

**WINDOWS NT[®]
SYSTEMS**

**PERFORMANCE
COMPUTING**

TOWER OF POWER



**Announcing the PowerRaid™ MAX . . .
the most cost effective, high performance fibre channel
storage system with speeds up to 100MB/sec**

A Solution in the Zzyzx SAN-Ready Product Line

- ▲ 100MB/sec FC-AL host or 80MB/sec Ultra2 LVD host channel
- ▲ 128MB battery backed cache
- ▲ Fault tolerant enclosures featuring hot swap power supplies, fans and disk carriers

- ▲ Flexible rack mount and tower configurations
- ▲ FC-AL solutions bundled with host bus adapters for all major operating systems
- ▲ Supports RAID levels 0, 1, 0+1, 3, 4 and 5
- ▲ Capacities up to 500GB

Available now featuring the latest generation disk drives by IBM

5893 Oberlin Drive
San Diego, CA 92121

Telephone: 800.286.7909
Fax: 858.558.8283

Zzyzx Peripherals, Inc.



Year 2000 Compliant

Visit our website at: www.zzyzx.com

(Zee-Zix): A road scratched into the middle of the desert in Southern California.

Joe Tier

Spinning An Outsourced Web

ISPs' systems-management procedures are an example of how to structure in-house methods, and what to expect if you outsource.

Overwhelmed by the time and resources needed for a successful Web undertaking, many companies are outsourcing the deployment and maintenance of their sites to managed hosting providers that provide security, backup, reporting, administration, and monitoring.

Many service providers use Intel-based servers running Windows NT Server, or SPARC-based servers running Solaris, and offer service-level agreements that guarantee fewer than five hours of unplanned downtime per year. Analysis of the reliability of these platforms by Digex Inc. shows that most failures are caused not by operating systems themselves, but by configuration and installation problems. Digex's studies also show that many seemingly OS-related performance problems and system crashes result from applications that contain custom code.

By following careful installation and configuration routines, examining

custom code, and configuring servers for fault tolerance and redundancy, many problems can be avoided.

cal data center and server-array architecture, and then discuss enterprise-management techniques.



In order to appreciate the scope of data-center management that is routinely required for large Web-hosting sites, we will explore a typi-

ARCHITECTURE

The diagram in Figure 1 (p. 30) illustrates network diversity and redundancy. Note that an optional firewall can be inserted into the configuration. A strong approach to high-availability servers begins with a multilayer server-array architecture. A strong Web-hosting company should have extensive experience in installing, configuring, and maintaining hardware and software. More significantly, the managed hosting provider should know how to provide standard guidelines for the physical architecture of a site.

A customer site might consist of a front-end Web server, an application server, and a database server. The site might include a firewall as a security measure but typically has no diversity or redundancy—each node is

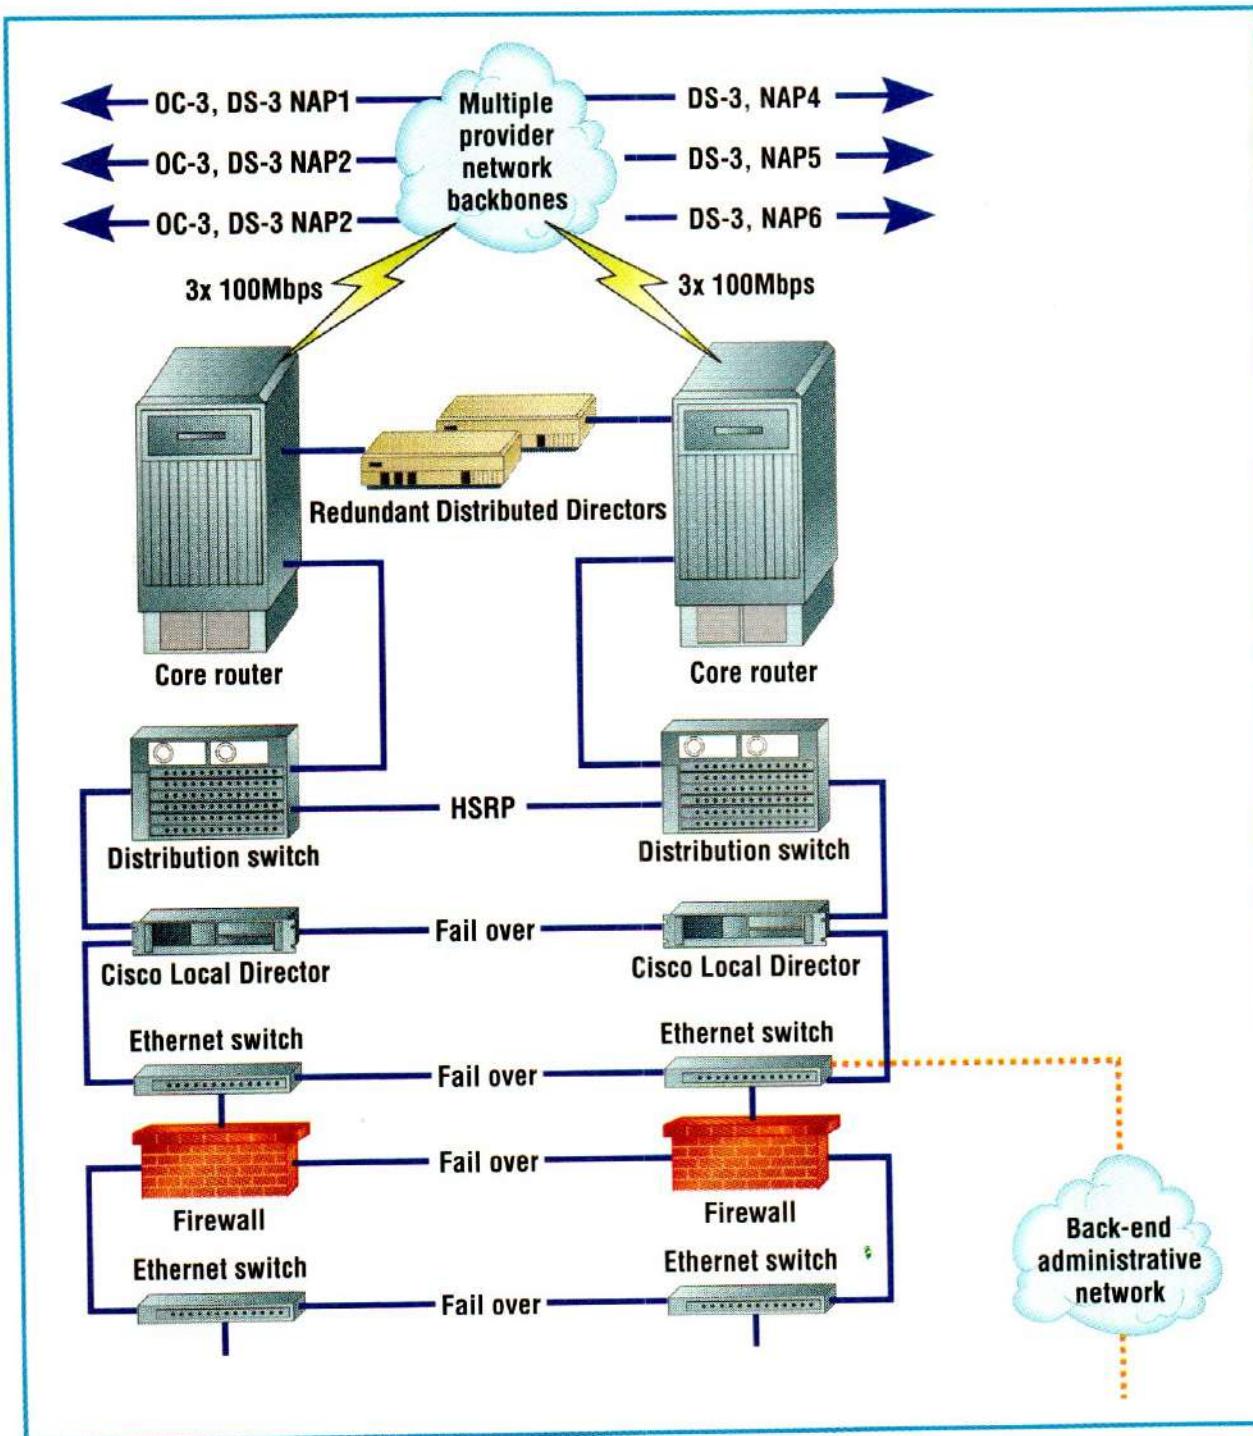
Dwight Bean

a single point of failure. Additionally, a T1 leased line is often used to connect directly from a site's physical location to its externally hosted site.

Upon further consultation with the managed hosting provider's engineers, the customer then decides to diversify and add redundancy as shown in Figure 2 (p. 32). At this point the site's servers and network devices are fully redundant and diverse, and a clustered database is utilized.

Similarly, a Windows NT site can be made redundant and diverse. The Windows NT Server infrastructure includes Domain Controllers and WINS servers in order to facilitate back-end NETBIOS connectivity between servers. Application servers can be made redundant using Microsoft's Windows Load Balancing Service (WLBS) or Microsoft Cluster Server.

Figure 1 Server network architecture



ENTERPRISE MANAGEMENT

A managed hosting provider's mission is twofold: manage hundreds of servers effectively and develop tools for managing thousands of servers. To accomplish this, enterprise-management tools are needed. There are two main aspects of enterprise management: remote server and configuration management, and monitoring. In order to resolve server/network issues, the monitoring function must detect the problem and trigger an alert. However, successful server management is proactive—server-management tools should be used before problems happen.

Remote Server/ Configuration Management

Detailed planning of how best to manage remote servers and their configuration is essential for both the ISP and the customer. Failure to define the

scope of system-management tasks will lead to disappointment with the services provided. The following are a few of the most important areas to consider.

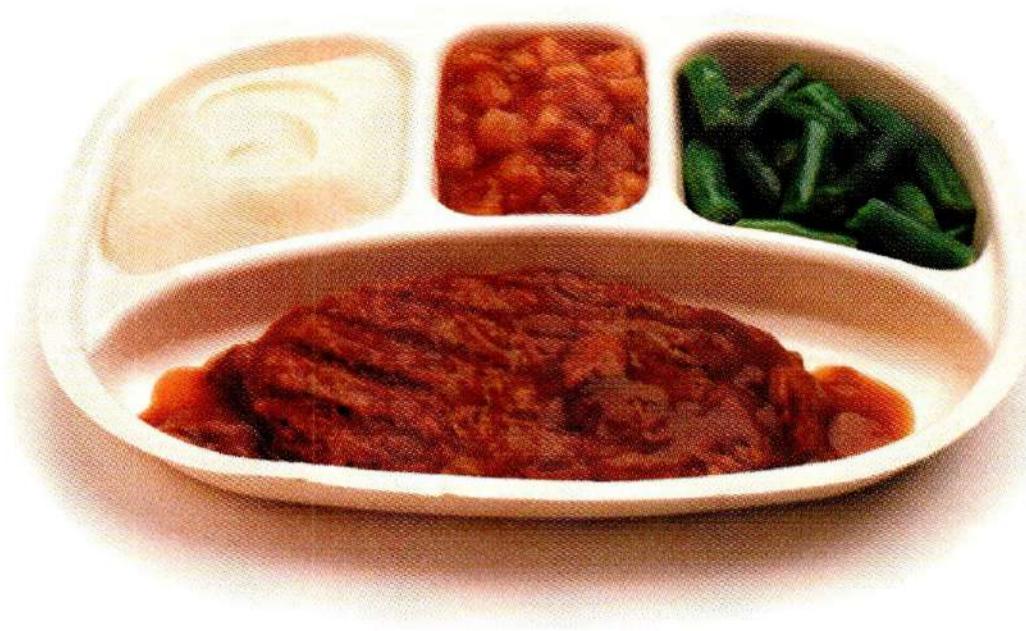
Software distribution. The managed hosting provider is required to maintain stable and security hardened OSs by continuously engineering the server build and distributing the latest patches, hot-fixes, and service packs to NT and Solaris servers. Tools that are commonly used for this function include Microsoft Systems Management Server for NT, and rdist for Solaris.

Asset/configuration/inventory management. A managed hosting provider is also required to execute asset, configuration, and inventory management through an electronic central nervous system.

Security management. Several aspects of security require monitoring. A security operations group should monitor newsgroups and security alert services for the latest OS vulnerabilities and patches. Real-time monitoring should occur in the areas of network (for example, a network management center), firewalls (a security management center), and servers (a server operations center). These groups use a combination of commercially available and internally developed applications to prevent system and network security breaches. For security reasons, discretion should be exercised in publicizing applications used for security monitoring.

Performance management. A managed hosting provider maintains stable and finely tuned OSs by using third-party applications as well. For example, NT performance can be managed via Performance Monitor and Compaq Insight Manager if Compaq is the NT hardware. Often, it is necessary to develop internal applications and scripts to manage NT and Solaris servers.

Network management. Network nodes are usually managed through telnet sessions. However, it is not uncommon to develop internal applications to provide Web-enabled (in-



SAME BASIC PARTS. DIFFERENCE IS, ONE WON'T COME BACK TO HAUNT YOU AT 3AM.



Parts are not parts. For the Linux OS, there's really only one company that puts all the right parts together in systems that are Linux specific from the chips on up. That's VA Linux Systems.

The industry's broadest selection of pre-configured, Linux tuned workstations, servers and even laptops are backed by a staff that includes some of the most renowned Linux developers in the world.

A recognized Linux leader, VA Linux Systems was first to ship Pentium III Xeon systems for Linux, first with two gig memory, first to develop one-terabyte storage systems and more.



An active member of the Linux community, VA Linux Systems operates and manages www.linux.com and provides nearly legendary telephone, web and 24/7 support.

Want to get the most out of your Linux OS? For ready-to-rock, off-the-shelf or fully customized Linux systems, no one provides faster, more reliable and more stable products than VA Linux Systems.

Why settle for anything less than all the right parts?



"The VarStation has been built from the ground up to run Linux and run it well. VA (Linux Systems) didn't scrimp on hardware and devoted a great deal of effort to ensure that the system is tuned to work as configured." InfoWorld

1.888.LINUX.4U www.valinux.com

Select No. 183 at www.performancecomputing.com/fyi/

© 1999 VA Linux Systems and its logo are trademarks of VA Linux Systems. Linux is a trademark of Linus Torvalds. All other trademarks are the property of their respective holders.

VA LINUX
Systems



Nothing runs Linux like VA Linux.

tranet) network-management tools. Commercially available applications such as What's Up Gold can be quickly deployed and configured as both a monitoring and management tool. What's Up Gold includes the ability to restart services and daemons.

General remote management. Solaris servers can be managed remotely through telnet, console port, and X Windows. Many NT functions can be managed remotely with Server Manager, SQL Enterprise Manager, and Microsoft Management Console (MMC), among others. There are many times, however, when NT console access is required. Although Microsoft's SMS remote-control and Symantec's pcAnywhere commercially available remote-control applications are valuable tools, they require network connectivity and can potentially mo-

nopolize CPU cycles. For live console access, many enterprises employ cascaded keyboard/video/mouse multiplexing switches, resulting in blocking factors of up to 60:1. (Blocking occurs when two users require access to servers in the same switch block at the same time.)

If access blocking becomes a problem, one solution is CCC Group's FreeVision X technology. FreeVision X lets live console access be transmitted as much as 1,000 feet from server to user.

Monitoring

Enterprise monitoring involves detecting each piece of the puzzle and then quickly putting those pieces together to see the big picture. For example, consider the following scenario:

9:06 A.M. Several monitored IP addresses begin to fail ping tests.

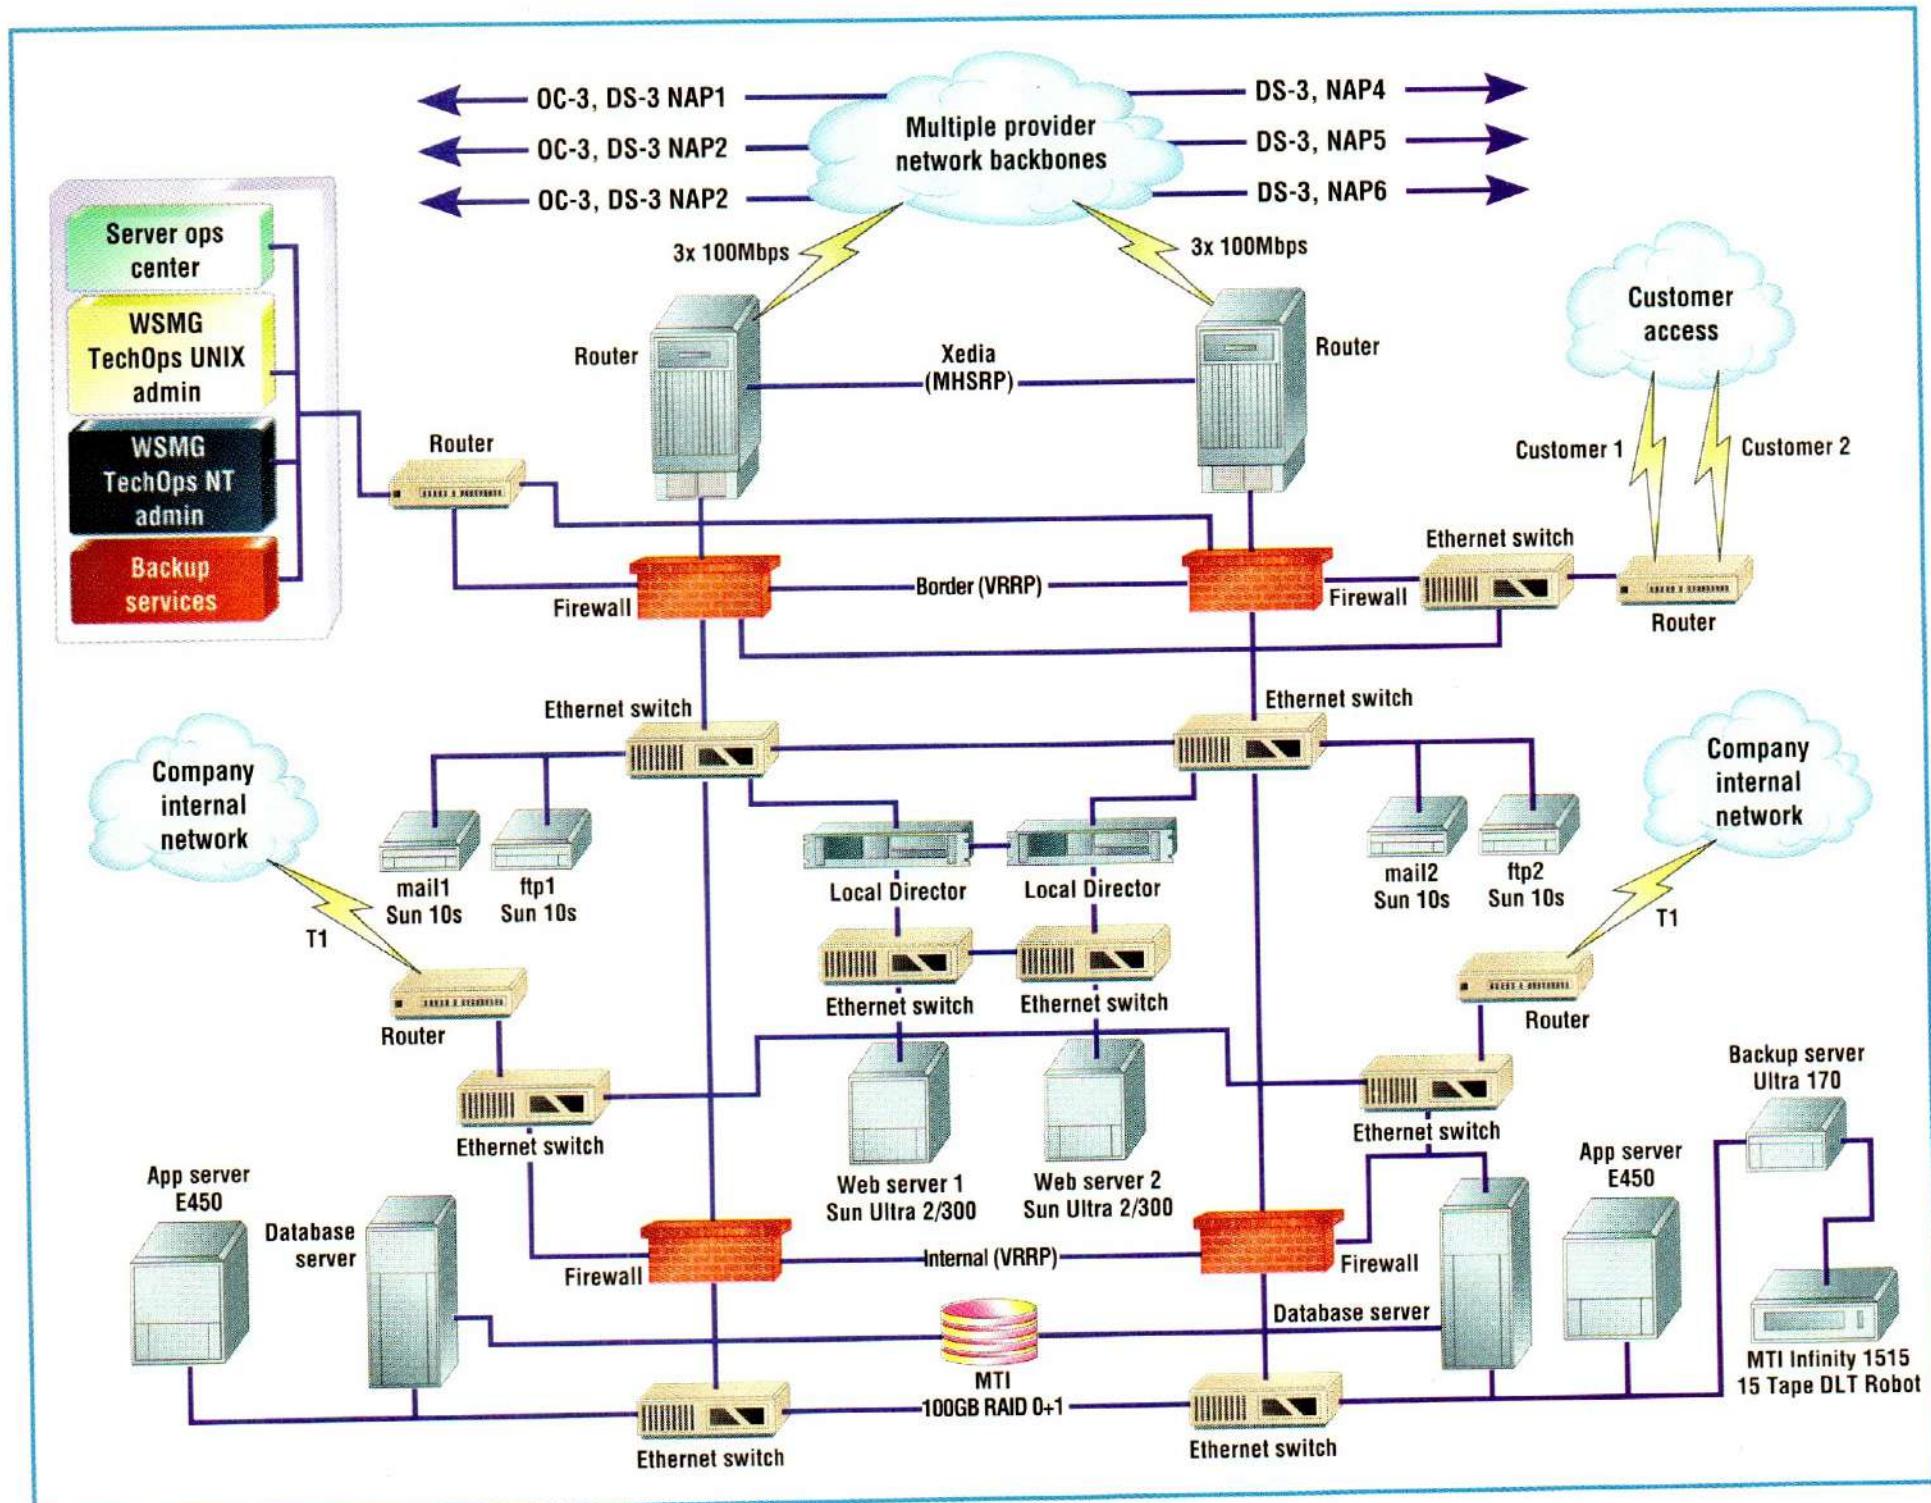
9:08 A.M. The <http://www.favorite-site.com> URL is detected as failing. Two more sites, <http://www.favorite2.com> and <http://www.favorite3.com> fail a few seconds later.

9:11 A.M. Two more sites fail. SOC technicians begin to notice a pattern. Has a switch failed?

9:13 A.M. An e-commerce site causes alerts that indicate the store is down.

9:15 A.M. Two more Web sites fail ping tests. Network and server operations engineers have joined in the troubleshooting process to determine that a network card is causing a "broadcast storm." The server with the errant network card is isolated from the network. This now causes another problem—the server is a key component in the e-commerce site. Fortunately, the customer has taken his technical account manager's ad-

Figure 2 High availability UNIX server array architecture



You'll be amazed that something this *broad and deep* can fit under one roof.



XML • Java •
C++ • UML • Enterprise Development •
Enterprise Application Integration • Enterprise JavaBeans • CORBA •
Component-based Development • Internet Development • Visual Modeling •
COM & COM+ • Windows 2000 • Testing • Security • Linux • User Interface

Software Development '99

Conference:
November 8-12

Exhibition:
November 9-11

Washington Convention Center, Washington, DC

The thing that makes Software Development '99 so exciting is the very same thing that makes it different from any other conference. SD'99 is the only major independent forum for development professionals offering the breadth of technologies and the depth of content you need to refine your skills for the next millennium.

This year, choose from more than 200 classes, with 120 new offerings that spotlight core languages and the hottest e-business technologies such as XML, Linux, UML, Java, Internet Development and Enterprise Application Integration. Then test-drive the latest tools from over 200 vendors on the Expo floor.

This November, come hear the visions of the industry's most influential leaders and share insights with your peers through social events and organized discussions. At SD'99, you'll join thousands of developers and their managers at the most critical event in the industry and be inspired to raise the roof with the next generation of software development.



www.sdexpo.com

call for registration information 800-441-8826 email us at sd99@mfi.com

vice and built a redundant configuration—the backup server takes over and the e-commerce site is operational again.

This scenario was resolved quickly by using data obtained from platform monitoring, application monitoring, hardware monitoring, general site monitoring, and reporting:

Platform monitoring. Platform monitoring includes the server hardware and general OS. For example, Compaq Insight Manager (CIM) uses SNMP traps to track the state of all NT server hardware components (if all of the NT servers in an enterprise run on Compaq hardware). These traps can be sent to a central group of CIM servers.

In the previous example, the network card's errors would be detected by CIM. The platform as a whole can be monitored through shareware applications such as NOCOL, which uses periodic ping tests.

Application Monitoring. NT and Solaris servers routinely run dozens of services and daemons, respectively. Upon deployment of a server, SOC configures What's Up Gold and Sitescope to monitor all services and daemons that are required for the server's operation. Upon failure of that service or daemon, attempts to restart are automatically executed in parallel with alerts sent to role e-mail and pager accounts.

General site monitoring. In the scenario above, an e-commerce Web

site only functions if all its components are running correctly. The best method of monitoring an e-commerce site is to determine which URLs require end-to-end functionality. For example, the "shopping cart," which is a standard feature on many e-commerce sites, is a great URL to monitor—the Web site service/daemon, e-commerce application, and database must all be functioning properly. Sitescope Omni-Vision can be configured to monitor these URLs.

Reporting. Enterprise monitoring requires a reporting function. There are two key areas in which reporting is necessary: event reconstruction and usage statistics. Event reconstruction includes postmortems and after-action reports. NT and Solaris both have in-

If you're thinking about outsourcing your

Web site, if you recognize the need to switch from technology management back to your core business, and if you're entering into e-commerce for the first time, there are some things you need to know.

Before you start, you need to determine which level of hosting services best fits your current, and more importantly, your future needs. Hosting-service providers fall into the following categories:

TYPES OF SERVICE

Shared hosting. Hardware and bandwidth is shared among a number of customers. Shared hosting is an economical entry point for companies that have relatively simple Web sites. The downside is that your site is likely to be limited in its functionality and more prone to slowing under heavy usage.

Collocated hosting. You own your own hardware, software, and network equipment housed at the service provider's facility and you are responsible for the installation, management, scalability, and security of your Web site. This can be an attractive alternative for Web-centric companies that have advanced in-house Internet expertise.

Dedicated hosting. The service provider supplies the hardware, software, network equipment, and support necessary to run your Web site. As Web sites are becoming more complex to meet the demands of increasing

e-commerce, many larger, technically astute companies are giving up trying to administer Internet technology in-house and are turning to dedicated hosting as an alternative to continual upgrades and migrations.

Application hosting. The service provider implements and configures a dedicated Web site along with managing certain Web-enabled business applications—for example, customer service, procurement, human-resource management, and sales-force automation—customized to your requirements. In general, application hosting offers benefits such as advanced application skills and lower cost of operation, but requires a certain amount of technology know-how on the part of the user.

Managed Web and application hosting. The service provider offers you both a managed Web site and application hosting, taking complete responsibility for the technical management of your e-commerce, while leaving you completely in control of its content. The most significant advantage of this approach is that it takes the entire onus of technology management off your desk and places it with your service provider.

If e-business is a critical element of your strategic plan, and you want to focus your company's time, resources, and money on your core business rather than on Internet-enabling technology, you're well advised to seek an e-partner whose capabilities rise to the managed Web- and application-hosting level.

—Melissa Hamilton-Sargeant

ternal logging capabilities. Cisco network equipment also has extensive logging capabilities. While review of individual server and Cisco equipment logs is often used, it is not a method that scales.

Cross-platform issues

When managed-hosting customers implement a cross-platform Web site, the major issues involve connectivity. Areas of connectivity primarily include file sharing and database access. For example, NT customers occasionally prefer a high-end Solaris server running Oracle. In these cases, Oracle SQLEasyNet is installed on the NT Web servers, and the connectivity issue is resolved. These issues may become further complicated by multiple security zones that are cordoned off with firewalls.

File sharing becomes an issue when a best-choice application is implemented on the "other" platform. For example, until recently it was perceived that the NT version of WebTrends Log Analyzer was the superior version of that company's software. Several all-Solaris customers added an NT machine to their configuration in order to make use of the WebTrends NT application. Again, the issue was connectivity, in particular, file sharing. In these situations, NFS applications for NT or FTP are used to transfer files.

TECHNICAL SUPPORT

As software tools, OSs, and applications get more complex by the month, a broader technical base is needed by today's IT professionals. In order to engineer, support, and maintain today's enterprise-management tools, expertise in four key areas is needed: OSs (and their associated hardware), networks (specifically TCP/IP), applications (including complex applications, databases, and programming languages), and generic troubleshooting skills.

Depending on their function within the managed hosting provider, differ-

ent groups require certain strengths. For example, consider what happens when an e-commerce Web site goes down. While troubleshooting the problem, an Internet System Administrator or engineer may need to understand the interdependencies of several Web site nodes: a Cisco Local Director, a Checkpoint firewall, a Solaris or NT Web server, a messaging server (either Sendmail or MS Exchange), an e-commerce application such as Broadvision or Microsoft Commerce Server, and a database (either Oracle

or MS SQL server). Resolving this problem requires a deep understanding of NT, Cisco equipment, and networking.

An SOC engineer needs to understand how to monitor each device or server. Understanding how all of the pieces fit together yields better monitoring solutions. For example, in order for a site's shopping-cart feature to be functional, the following nodes must be functioning: the Cisco Local Director, checkpoint firewall, Web server daemon or service, SQL database (a

Enterprise-management tools

- **NOCOL** A shareware port-monitoring tool that can execute success/failure tests against TCP/IP ports.
- **What's Up Gold (Ipswitch Software)** An enterprise-monitoring tool that provides real-time views of network and server status and presents the data in a hierarchical display (<http://www.ipswitch.com>).
- **OmniVision (Systar Software)** An enterprise-monitoring tool that provides real-time views of network and server status and presents the data in a business-format display. OmniVision also includes reporting and data-mining tools (<http://www.systar.com>).
- **Sitescope (Freshwater Software)** An enterprise tool that monitors several server functions and presents the data in a "dashboard" display (<http://www.freshwater.com>).
- **Compaq Insight Manager** An SNMP-based enterprise tool that monitors Compaq hardware components (<http://www.compaq.com/>).
- **NetCool (Micromuse)** An enterprise tool that provides real-time network monitoring. It also includes an open architecture and a library of off-the-shelf software modules that collect event information from over 100 popular network environments (<http://www.netcool.com>).
- **Systems Management Server (Microsoft)** An enterprise-management tool that includes software distribution, software inventory, and help desk functionality (<http://www.microsoft.com>).
- **WebTrends Log Analyzer (WebTrends)** One of the most popular and powerful Web-site-traffic-analysis packages on the market today, and one that quickly processes usage-log files (<http://www.webtrends.com>).
- **Solaris RDIST (Sun Microsystems)** An enterprise-management tool that includes software distribution, software inventory, and help-desk functionality (<http://www.sun.com>).
- **FreeVision X (CCC Group)** FreeVision X is a video-multiplexing technology that lets, for example, 32 users access any of 1,024 servers right from their desktop. FreeVision facilitates an NT "lights out" environment (<http://www.cccgroup.co.uk>).

For the type of management and monitoring required by most hosting service companies, an off-the-shelf software package or suite of applications that's ready to go does not exist. For this reason, many providers are currently developing their own solutions, which typically combine best-of-breed commercial applications with internally developed scripts, rulesets, and applications to get a solution that works best for the service provider and its customers.

recent set of consistency checks should have been executed), and Microsoft Commerce Server service. Although each of these services can be independently monitored (and often are), end-to-end monitoring must occur. In this case, the SOC engineer should work with the customer to determine URLs within the Web site that test the successful integration of the

site's individual components. Generating other monitoring solutions can also involve the rapid development of scripts. Several of the monitoring tools mentioned here incorporate the ability to call custom applications and scripts.

A Technical Account Manager (TAM) is required to understand all of these disciplines at a higher

level—not at the day-to-day nuts-and-bolts level, but at a level in which system architectures can be documented and discussed. Further, TAMs will often be required to identify potential problems before the fact. For example, when asked to make a specific change as requested by the customer, the TAM should use his or her expertise in the core areas to determine the viability of a requested action. If the requested action doesn't make sense or looks like a potential problem, the TAM should go back to the customer and determine the overall goal of the requested action. The TAM can then recommend a consistent set of actions that will meet the desired goal rather than risk precipitating a potential failure.

When a managed-hosting provider develops and exercises these areas of expertise throughout the company, there is a true value-added service that justifies the word "managed."

SUMMARY

Launching and maintaining high-availability Web sites is not as easy as it sounds. The Internet is clearly a leading-edge (if not "bleeding-edge") technology-driven environment. For this reason alone, the stakes are higher just in contracting for over-99-percent-upptime service level agreements. There are many times when managed-hosting providers push the envelope of the vendors whose products they use. A week rarely passes where the typical service provider isn't returning to a vendor with solid evidence of new issues not previously encountered in their software or hardware. The rigorous and perpetual pounding of users and network traffic provides the ultimate testing process for solid applications, OSs, and network/server hardware. This in turn reinforces the need for constantly improving enterprise management and monitoring tools and the staff to develop, configure, and maintain those tools.



CLASSIC AFFORDABILITY

J. River's ICE.TCP is designed to be an easy-to-use, high performance Windows to UNIX connectivity suite.

Over 1,500,000 people enjoy its unparalleled depth of terminal emulation, file transfer and printing.

You'll be pleasantly surprised to find how affordable this standard is and why customers with 1 to 10,000 users choose J. River's ICE.TCP for price and performance.

We invite you to visit www.jriver.com to obtain a trial version.

J. River


www.jriver.com

Affordable Network Gear
 612-677-8200

Select No. 22 at www.performancecomputing.com/fyi/

Get every issue of *Sys Admin* ever published— on one CD-ROM!

NEW
VERSION 4

- Seven years of valuable tips, tools and techniques—at your fingertips!
- Easy access to practical time-saving articles and code.
- Best of all—you can use it on any system with a CD-ROM drive!

Every issue of *Sys Admin* from 1992 through 1998 on one convenient CD-ROM!

Since 1992, *Sys Admin* has provided UNIX systems administrators with the best practical technical information for keeping UNIX systems up and running. Included are articles by Bjorn Satdeva, Arthur Donkers, and Elizabeth Zinkann.

Put seven years' worth of *Sys Admin* at your fingertips — no more fumbling through back issues, searching for that practical article on UNIX accounting facilities or the code for configuring FTP on your SVR4 system.

Key Topics include:

- backup
- networks
- Web administration
- ...and much more.
- security
- performance
- UNIX/DOS/NT connectivity

It's the ultimate resource for UNIX systems administrators!

both HTML & ASCII formats



**Sys
Admin**
CD-ROM

1992-1998 Call 1-800-444-4881 Now!

SA

Sys Admin CD-ROM 1992-1998

New Version 4

7 years of *Sys Admin* articles & code, including:

- 58 issues from the authority on UNIX systems administration.
- Hundreds of valuable administration tools, tips & tricks!
- Complete code listings.

ONLY \$49.95
(plus shipping & handling)

UPGRADE from version 3 \$24.95
(plus shipping & handling)
s&h \$3.50 USA/Canada
\$13.50 all others

Phone: 800-444-4881, order code SACD4
785-841-1631

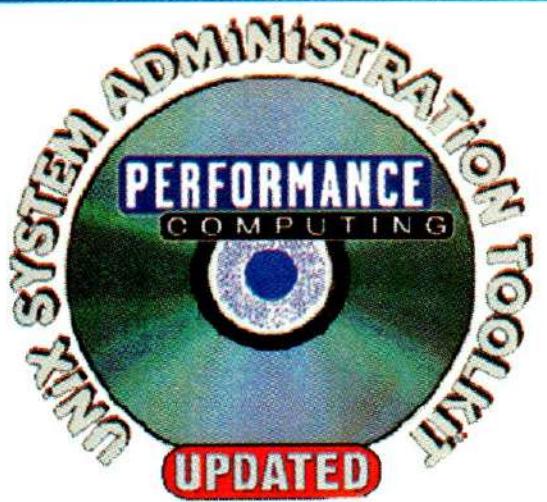
Fax: 785-841-2624

Email: orders@mfi.com

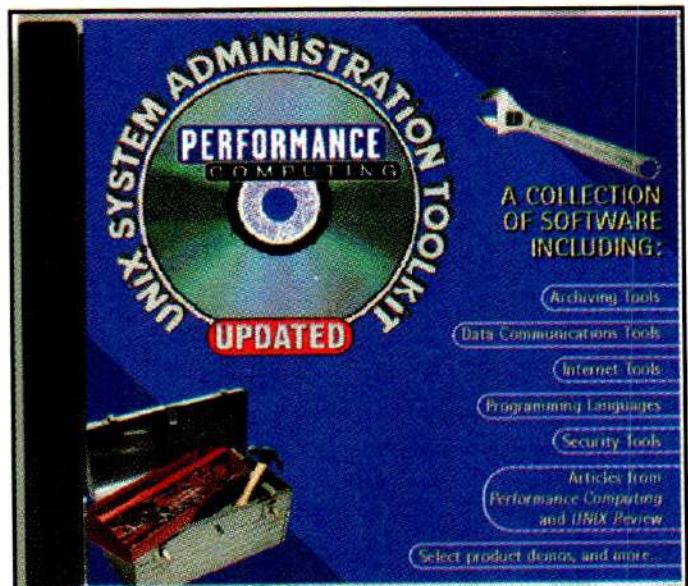
Mail: 1601 West 23rd St., Ste. 200
Lawrence, KS 66046, USA

Web: <http://www.samag.com>

Have All The Tools You Need



Order the *Performance Computing*
UNIX System Administration Toolkit CD-ROM



This low-cost comprehensive CD-ROM is packed with tools for the UNIX system administrator. This updated edition includes the latest versions of these tools, thus saving you hours of search and download time. It also includes source code, select software demos from leading solutions providers, and a library of reference material from past issues of *UNIX Review* and *Performance Computing*.

Check out these features*

- Programming Tools and Languages
- Data Communication Tools
- Security Tools
- Performance Monitoring Tools
- Web Server Software
- Web Browsers

*For a complete list of included applications go to www.performance-computing.com

Plus...

Select Solutions-oriented Product Demos and a Comprehensive Library of *UNIX Review* and *Performance Computing* Editorial Content Including an Archive of System Management Feature Articles "Perl Advisor" and "Daemons & Dragons" Columns

ALL FOR JUST \$39.95!

ORDER FORM

YES! Rush me ____ copies
of the *UNIX System
Administration Toolkit CD-ROM*
at \$39.95 each!

Quantity	Price	Total
	\$39.95	\$
		Subtotal
		Shipping/ Handling
		Sales Tax
		TOTAL \$

**ORDER
TODAY!**

Complete this form and return to: *UNIX System Administration Toolkit CD-ROM*, 1601 West 23rd Street, Suite 200, Lawrence, KS 66046. Telephone: 800/444-4881 or 913/841-1631; FAX 913/841-2624; e-mail orders@mfi.com

Name	Company Name		
Address			
City	State	Zip	Country
Phone	Fax		
<input type="checkbox"/> Check	<input type="checkbox"/> VISA	<input type="checkbox"/> MasterCard	<input type="checkbox"/> American Express
Credit Card#	Expiration Date		
Signature			

Sales Tax: Residents of NY-8%, IL-6.25%, CA-7.5%, TX-7.5%, KS-6.9%, GA-5%

Shipping & Handling: \$4 for U.S. and Canada orders; \$10 for international orders

Ralph Barker

IBM'S RS/6000 HA-H70 CLUSTER SERVER

You may not want to wear it on your wrist, but this 850-pound cluster keeps on ticking.

The Internet and the Web are credited with effecting profound changes on various aspects of our lives, including our perception of time. Development and business trends now take place during "Internet time," meaning that the pace of whatever is being discussed is now vastly accelerated. A side effect of this acceleration is that "upness," otherwise known as availability, is now a critical concern in almost all businesses. If your systems are not up, you are losing business, and both your profitability and stock price are likely to be affected. While this has long been the case in data-center operations, where financial and other forms of business-critical computing take place, the Web has broadly extended this concern.

The solution to the availability challenge that surfaced a decade or so ago, when minicomputers started encroaching on what had previously been mainframe territory, was hardware redundancy. Tandem Computer (now part of Compaq), for example, became successful by offering systems in which almost every critical part or subassembly had a hot spare ready to take over in an instant. A failed hard-



ware component no longer meant a crashed system, so the computer-based business operation was able to continue without interruption. A great design, but one with the obvious downside of high cost. Doubling the part count not only doubled the cost, but also added a level of complexity to the system's design, and required additional circuitry to make the failover possible. The current thinking is that every computing environment might require a similar level of near-absolute availability.

The cost/benefit ratio of complete redundancy in systems, however, prompted a different approach—that of highly available (HA) systems. Here, a few seconds or a few minutes of downtime, along with some level of inconvenience to users, is an acceptable tradeoff to the cost of fully-redundant system design. The definition of HA varies, of course, between both manufacturers and users. In fact, most vendors of HA systems offer varying degrees of HA through options to their basic HA offering. Thus, the level of HA can be tailored to your specific availability requirements. This month's Tested Mettle examines one such offering: IBM's HA configuration of its H70 servers, which are part of the broader RS/6000 line.

OF NOTE

High availability is probably best thought of in terms of a layered, or almost hierarchical, architecture. Depending on the requirements, an HA system may involve both a cluster of servers and load-balancing components that might be external to the cluster. Within the cluster itself, most HA configurations, including the HA-H70, offer considerable flexibility in the setup, and can be tailored to the user's specific HA objectives. A common configuration is one in which the servers making up the cluster are "aware" of each other, share certain resources, have a mechanism for detecting a failure within the cluster, and can perform work indepen-

dently of each other. This is how we configured the HA-H70 during our testing, but other HA configurations can be structured to have one or more systems in the cluster acting as hot standbys, ready to take over operations, but otherwise sitting idle while they wait for another system in the cluster to fail.

An aspect of HA that should be recognized early is that HA is not solely a hardware-based solution. Depending on the nature of the applications running on the servers in the cluster, the applications themselves may need to be "cluster-aware." For example, a sophisticated database application might need to have cluster software components that will allow the application to roll back transactions that were started on the failed server, or allow the transaction to be completed once the user is moved to the server that has taken over operations for the failed machine. Thus, the orchestration of the failover process must take place at the OS level among the constituent systems, and may be further complicated by associated operations that need to take place at the application level, as well.

The basic concept underlying IBM's HA implementation is that the nodes in the cluster have defined roles, and share various defined cluster resources. The roles played by the cluster nodes vary based on the style of cluster configuration you select. In the typical configuration, one of the nodes acts as the cluster manager, in addition to doing its own share of work, and "owns" certain defined cluster resources. When a failure within the cluster is detected, a previously scripted sequence of events unfolds through the cluster software, and transfers ownership and control of cluster resources to the other node. Work progresses, though perhaps more slowly, until the failed node is repaired and rejoins the cluster. This design provides what IBM describes as "nearly continuous availability" for mission-critical applications.

The software at the core of the HA-H70 is called High Availability



Cluster Multiprocessing for AIX (HACMP), and the version installed on our review system was 4.3. HACMP works with the underlying AIX operating system, and standard system and network components, to provide the configured HA features. In the configuration of our review system, redundant data paths were provided by dual NICs and SSA interfaces in each of the cluster's two servers. Network-status monitoring and inter-node communication (essentially the heartbeat of the cluster) occurs over the standard network using relatively low-overhead protocols. A separate, high-speed network typical of a ccNUMA cluster is not required, although some similarities between the architectures do exist.

HACMP configuration and management can be accomplished through several avenues. HACMP includes both character and graphical enhancements to AIX's Systems Management Interface Tool (SMIT). Additionally, the software includes a series of command-driven utilities, and graphical utilities for display under the Common Desktop Environment (CDE) GUI, as part of HACMP's Visual System Manager package. The software's Cluster Single Point of Control facility lets you perform common cluster-management functions from either node in a two-node cluster, regardless of which node currently "owns" cluster resources. HACMP also includes its own Management Information Base (MIB), so it can be integrated with various SNMP-based high-level management systems. Thus, HACMP provides considerable flexibility with respect to its configuration and management.



*So many products...
so little time.*

*Performance Computing's
Online Buyer's Guide*

Expanded format for 1999,
with 350 vendors and almost 1000
products—updated monthly

Search by keyword, company name, product
name, product category, or product type

*We're
making the search a
little easier.*

www.performance-computing.com

As to the hardware itself, the HA-H70 comprises two RS/6000 Model H70 servers and an IBM 7133 Model D40 Serial Disk System, all mounted in an IBM 7014 Model S00 rack. Each H70 server can include up to four 340MHz RS64 II microprocessors, each with 4MB of L2 cache. A minimum of 512MB RAM is required, and can be expanded to 8GB per server. Each server also includes dual Ultra2 SCSI controllers, dual boot disks, dual advanced serial RAID PCI SSA disk controllers, and dual LAN adapters (various network topologies are supported). As might be expected, the D40 disk system also provides redundant data paths.

OPERATION

Installation of the HA-H70 requires the level of preplanning associated with any system of this sophistication. As usual, you will want to consider power and thermal requirements in terms of the location in which you intend to install the system. The S00 rack includes a power-distribution box and a single 220V power cable with a standard three-prong, twist-lock connector. Each fully configured H70 drawer weighs about 195 pounds, and pulls about 750W of power. The 7133-D40 disk array weighs another 118 pounds and uses about 430 watts, while the S00 rack itself weighs almost 350 pounds. Thus, the HA-H70 unit weighs in at close to 850 pounds, and needs almost 2,000W on a 220V line.

Before installing the HA-H70, you will want to review the HACMP documentation, and it is wise to request copies of the manuals well in advance of the arrival of the system. In particular, look for the booklets "Concepts and Facilities," "Planning Guide," and "Administration Guide." These manuals provide a detailed discussion of HACMP concepts, and include worksheets that will be helpful in determining your HA configuration. There is also a hefty "Installation and Service Guide" manual for the model H series servers that you will want to become

familiar with. Much of IBM's documentation is also available online at http://www.rs6000.ibm.com/resource/hardware_docs/index.html. At review time, no H70 technical documentation was available there.

Each of the H70 drawers in the cluster provides considerable room for expansion. Each unit has eight PCI slots; four 32-bit, and four 64-bit. Each unit comes standard with diskette and CD-ROM drives, and has a third media bay that can be used for tape, CD-ROM, or hard disk. Twelve additional bays are provided for 4.5GB (or larger) hot-swappable hard disks, with a maximum disk capacity of 254.8GB in an H70 drawer. Three serial ports are provided on the rear of each unit, along with one parallel port, an external SCSI connector, an integrated 10/100Base-T Ethernet port, an AUI Ethernet connec-

tor, and two optional SSA connectors. Dual power supplies are accessible from the rear of the unit, along with the CPU and I/O board fans, both of which are hot-swappable. Two additional fan units are mounted at the front of the system: one for removable media drives, and one for the hard disks.

The 7133-D40 disk drawer provides the additional shared Serial Storage Architecture (SSA) for the cluster. The D40 supports a mix of disk sizes, including 4.5GB, 9.1GB, 18.2GB, and 36.4GB SSA drives. The unit accommodates up to 16 drives, and IBM recommends the drives be connected in two SSA loops for optimum performance.

Although the HA-H70 includes two H70 servers in the HA cluster, the HACMP software can handle up to four nodes, or up to 32 nodes with the op-

IBM RS/6000 HA-H70

IBM Corp., RS/6000 Division
P.O. Box 100, Somers 3, Rte. 100
Somers, NY 10589
800-IBM-CALL
<http://www.ibm.com>

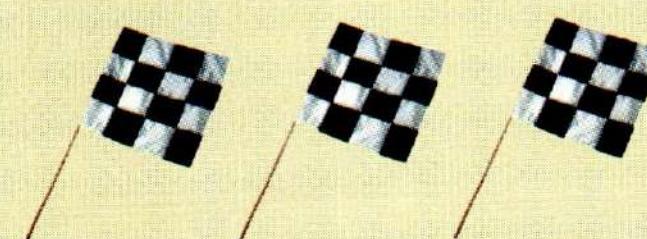
TESTED CONFIGURATION: The HA-H70 comprises two IBM RS/6000 Enterprise Servers Model H70, and an IBM 7133-D40 SSA disk system with eight drives (4GB), all housed in an S00 rack. Each H70 server included four 340MHz RS64 II microprocessors, each with 4MB of level 2 cache; 2GB RAM, six internal hard disks (4GB); Dual Integrated UltraSCSI-2 controllers, Dual 4.5GB UltraSCSI IPL boot disks; Dual advanced serial RAID PCI SSA disk adapters; Dual 10/100 Mbps Ethernet LAN adapters, IBM PCI Token Ring; Integrated 10/100 Ethernet PCI, Ethernet 10Base-2 Transceiver; Redundant AC power supply; CD-ROM drive, 1.44MB 3.5-inch diskette drive, 8mm

tape drive; HACMP Version 4.3; and AIX Version 4.3.2 (1-2 user).

PRICE: \$224,255

OPTIONS: Various add-ons to HACMP, and extensive hardware expansion and networking options. Starting price for the HA-H70 with two one-way H70s, 512MB RAM per server, and eight 4.5GB disks in the 7133-D40 is \$108,700. The minimum configuration price for a standalone H70 server with 128MB RAM and one 4.5GB disk is \$26,900.

EVALUATION: The 340MHz RS64 II is not a barn-burner CPU, but provides respectable performance on industry-standard benchmarks. The H70 servers provide outstanding expansion capability, however, and the HACMP software is a robust part of IBM's midrange HA solution. The IBM RS/6000 HA-H70 Cluster Server earns a good overall rating, and the associated three *Performance Computing* flags.



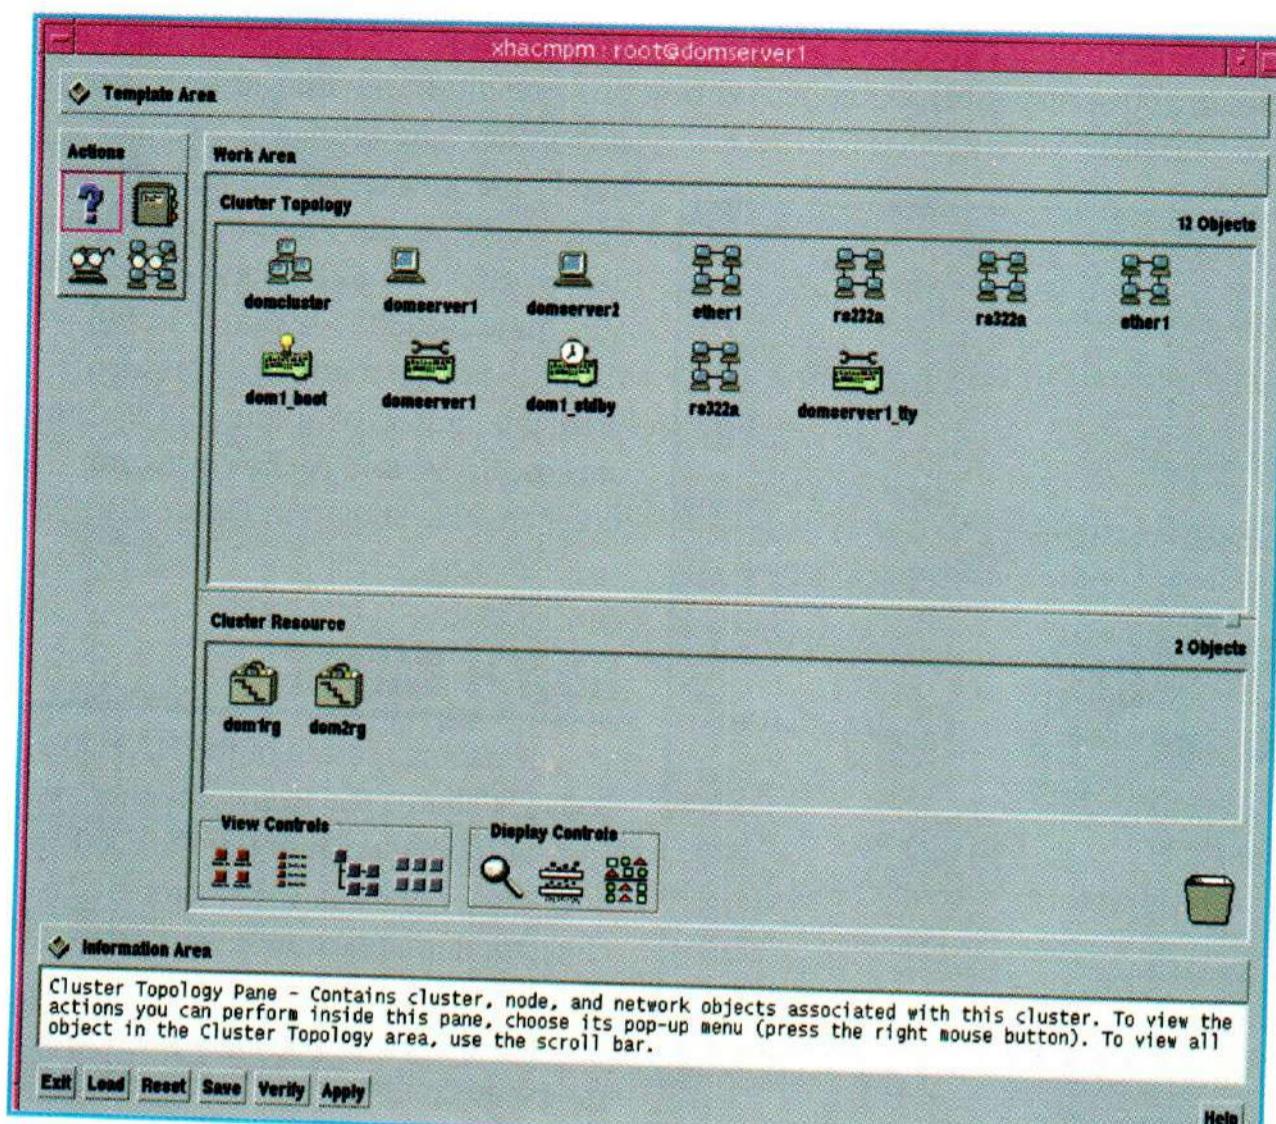


Figure 1 The HACMP interface

tional Enhanced Scalability feature. Thus, a cluster can comprise up to 128 CPUs.

The HACMP software is layered on top of AIX, so most standard AIX commands and administrative utilities function normally. HACMP commands are added to both the character-based and graphical versions of SMIT, or can be run from the command line, as mentioned earlier. Figure 1 shows the graphical HACMP interface provided by the xhacmpm utility.

Four HA configurations are supported by the software:

- **Idle standby**, in which one otherwise idle node in the cluster functions as the standby replacement for a failed node
- **Rotating standby**, where a predefined or contention takeover sequence provides the standby capability
- **Mutual takeover**, where all nodes in the cluster back up each other while sharing the application load
- **Concurrent access**, where working nodes share the same workload and data simultaneously

With a concurrent-access configuration, the optional LoadLeveler for AIX software can allocate user logins across the nodes in the cluster based on utilization levels, thereby maximizing overall cluster performance.

PERFORMANCE

As systems become more complex, so does the mix of criteria that should be considered when assessing and comparing the machine. Performance at the high end is not just a matter of having the fastest CPU. A combination of processing power, system bandwidth and latency, and architectural issues affect the system's overall performance rating. This is particularly true when looking at HA systems, where the complex combination of hardware and software features determines a system's suitability to a particular task in a specific environment. A mix of features that may be perfect for one location may not be suitable for another environment. Thus, we felt that a broader view of the HA-H70's performance was called for, rather than just looking at SPEC benchmark performance.

On the HA side of the HA-H70 equation, performance considerations are particularly site-dependent. HA is not instantaneous—failover can range from a few seconds (nine to ten seconds minimum on the HA-H70) to several minutes, depending on the complexity of the configuration, the resources shared within the cluster, and the application failover processes that may need to take place. While it would be convenient to be able to distill the HA-H70's HA performance down to a single number that could be easily compared to other systems, that is not possible. A comparative analysis done by D. H. Brown Associates Inc. is available on IBM's Web site (<http://www.austin.ibm.com/resource/consult/dhbrown/rasc.html>) and is worth a bit of study. The study examines 11 categories of HA functionality, and compares various high-end systems, applying certain weighting factors to each of the categories. The study also cautions that the weighting factors used in the comparisons may not be applicable in all situations, so you will have to do a bit of digging for the information that is most useful to your environment. Bear in mind when reading the report that it examines an earlier version of HACMP running on a different RS/6000 system, an S70. The information is useful, nonetheless, and much of it can be applied to the HA-H70 with a little updating. Benchmark results are far easier to deal with.

One of the target markets for the HA-H70 is Web-driven e-commerce. With respect to Web performance, the system's SPECweb96 results are helpful. A four-way H70 server scored 11,774 HTTP operations/sec in that test. For comparison, consider the Dell PowerEdge 6350/500 running four 500MHz Pentium III Xeon CPUs (13,984 ops/sec), the Hewlett-Packard HP9000 N4000, running four 440MHz PA8500s (13,051 ops/sec), and an older Sun Enterprise 450 running four 400MHz UltraSPARC II CPUs (9,115 ops/sec).

For conventional CPU performance, we used our normal battery of SPEC95 benchmarks. Our SPEC95 runs on one of the nodes in the HA-H70 cluster

produced scores of 648 for SPECfp_rate95 and 558 for SPECint_rate95. As shown in Figure 2, the H70 scored

higher on floating point than SGI's Origin 200 server (614), equipped with four 225MHz Mips R10000 CPUs and

256MB RAM, but was lower than the SGI's result on the integer operations (609). Sun's Enterprise 450 scored higher on both tests (758 and 609 respectively), as did the Compaq Alpha-Server 4100 5/600 (858 and 657) and HP's new N4000 (1,495 and 1,209). All of the comparison systems also were configured with four CPUs. Chip speeds, RAM, and cache configurations are also shown in Figure 2.

HACMP features

Some of the other features of High Availability Cluster Multiprocessing for AIX (HACMP), as explained in the IBM documentation, include:

HORIZONTAL SCALABILITY

- Boosts overall performance and capacity by sharing disk and/or processor resources of clustered systems, thereby spreading applications across RS/6000 servers
- Provides scalable growth with reduced reinvestment and increased system availability
- Enables mix of uniprocessor and multiprocessor nodes for application performance and disk sharing

SUPPORT FOR MULTIPLE AVAILABILITY CONFIGURATIONS

- Allows tailoring of HACMP to your environment for a flexible solution that can change with your business
- Allows growth for availability in either standby or mutual takeover modes in 2- to 32-node clusters

CLUSTER MANAGER

- Advises of HACMP cluster configuration changes
- Allows access to all shared data within the cluster even if a system in the cluster fails

TAKEOVER SCRIPTS

- Offer several different ways to customize takeover actions in a cluster

SUPPORT FOR CLUSTER MANAGEMENT SMIT AND VSM FACILITIES

- Provides common interfaces (such as a drag-and-drop GUI) to make it easier to install and configure highly available cluster systems as well as maintain them on the network
- Allows cluster management from a single system console for the cluster
- Provides comprehensive highly available monitoring utilities to manage and tune clusters
- Enables visualization of relationships between cluster hardware and software resources

CLUSTER SINGLE-POINT-OF-CONTROL

- Enhances system management by helping to reduce system-administration errors and enabling system tasks to be performed once for all cluster nodes
- Provides a single-server image so the cluster looks like a single system for reduced system administration

DISTRIBUTED LOCK MANAGER

- Enables additional application-level facilities that can help coordinate concurrent executions, such as resource sharing on different platforms
- Lets applications simultaneously coordinate execution on complex clusters
- Provides data integrity when sharing common data in a concurrent access cluster

CLUSTER NODE SNAPSHOT UTILITY

- Saves configuration and configuration changes to facilitate cloning of additional clusters
- Allows maintenance of multiple cluster configurations

DYNAMIC RECONFIGURATION

- Enables continuous system maintenance without disruption
- Allows cluster resources to be added or removed without disruption

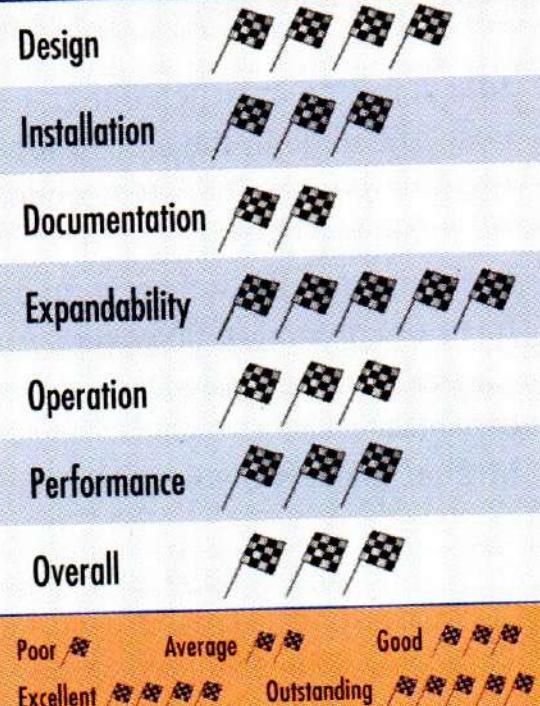
HOW IT RATES

The cabinet design of the HA-H70's S00 rack shares features with other systems at the high end of the RS/6000 line—matte black in color, strongly architectural in appearance, and generally distinctive. The system engineering inside the cabinet provides ample room for expansion boards, RAM, and disk storage within each of the H70 units. Additionally, the system's design forms a good foundation for HA features, and we liked the structured robustness of the HA features within the HACMP software. We rate design as excellent for the HA-H70, and confer the corresponding four *Performance Computing* flags.

Installation of the HA-H70 was consistent with that of other clustered systems we have worked with, although configuration of the

THE FINISH LINE

IBM's RS/6000 HA-H70



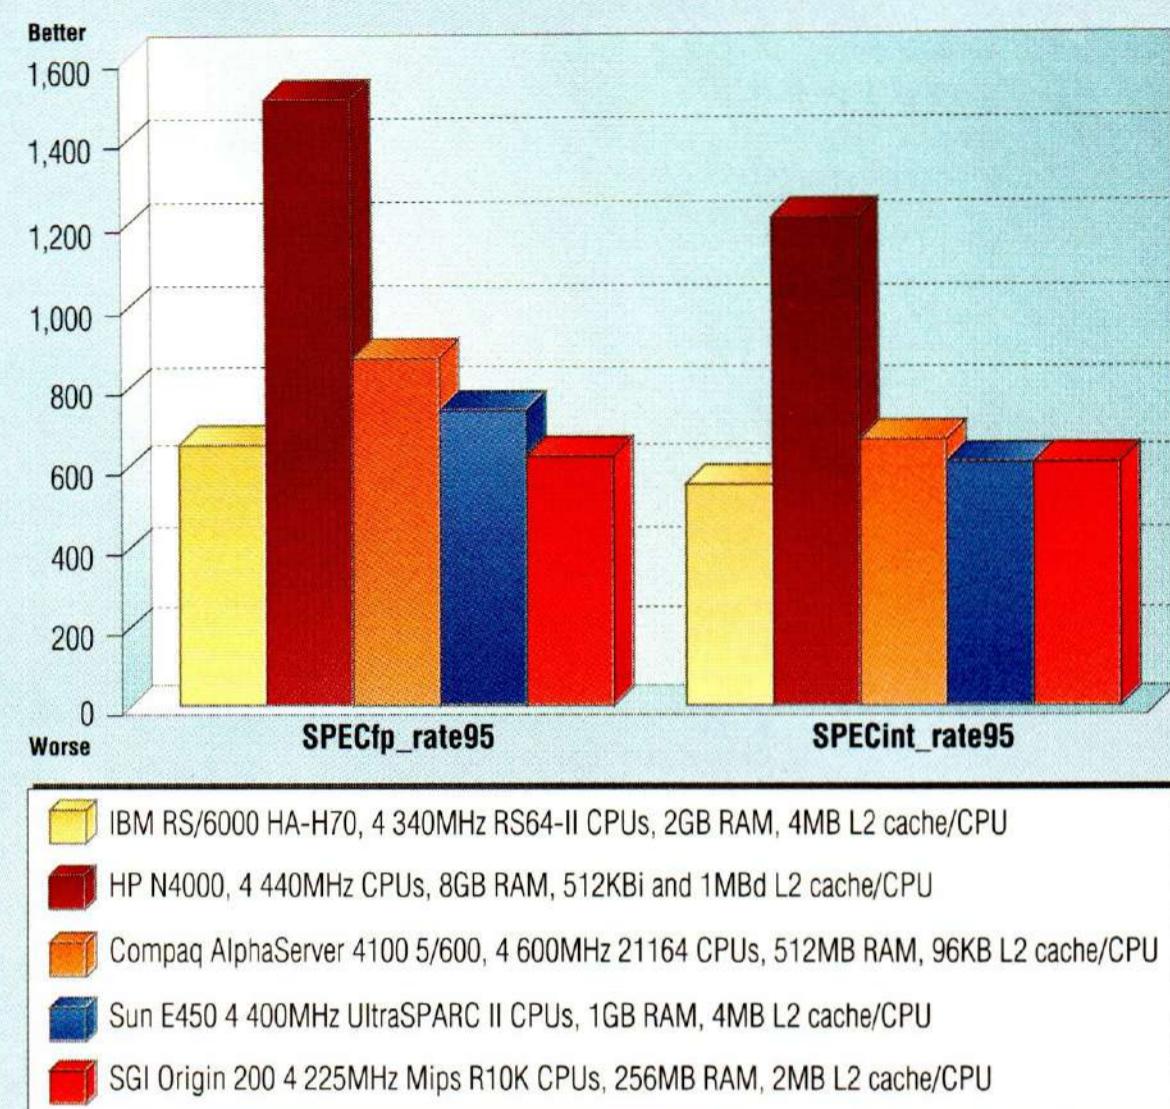
HACMP software was more labor-intensive than we expected. IBM does provide software engineering support, as do most manufacturers of high-end systems. Thus, if you are new to HA, taking advantage of that extra level of support can be beneficial. We encountered some trouble in getting the appropriate pre-installation documentation, as the HA-H70 manuals were not available online at the time of our review. We rate installation of the HA-H70, including the extra engineering support, as good—three flags.

Printed documentation for the HA-H70 hardware and the HACMP software is generally well-written and accurate. The HACMP manuals could be improved with better introductory sections, however. The manuals tend to delve into detailed descriptions of how the tree limbs work without first explaining what either the trees or the forest are all about, and so working your way through the material can be a bit tedious. You will also want to install all of the available online documentation. It is not installed by default at the factory, even for a system of this size and complexity, and having it available online locally is far more convenient than dealing with delays in Internet flow. We rate documentation for the HA-H70 as average—two flags.

Expandability within each of the H70 server units, and the S00 cluster cabinet, is excellent. We were impressed by the 8GB RAM capacity of these relatively small units, and the H70 has more internal disk capacity than any system this size we have worked with. Although the SSA storage arrays might be considered dead-end technology, because SSA has largely been supplanted by Fibre Channel, the SSA architecture fits nicely with the cluster requirements of the HA-H70, and provides excellent storage density and bandwidth. The HA-H70 earns an outstanding in the category of expandability, a full five-flag rating.

The HA-H70's operation displayed robustness during our testing. We had no hardware problems during our review of the system, and AIX 4.3.2 pro-

Figure 2 SPECrate95 benchmark results



vided a solid foundation for the clustering software. HACMP detected our artificially introduced hardware glitches quickly, and failovers took place as expected (within about 15 to 20 seconds on our relatively simple configuration). It should be noted, however, that our S00 came with but a single power cable, meaning that if one electrical circuit goes down, so does the whole system, clustered or not. We rate the HA-H70's operation as good, another three flags.

Performance of the HA-H70 on industry-standard benchmarks was respectable, but not stellar. The HA-H70's SPECint_rate95 result was close to that of competing midrange systems from Compaq, SGI, and Sun, but trailed Compaq and Sun on floating point. Further, the HA-H70 was substantially behind the scores of the HP N4000, but so were the other systems used for comparisons. The HA-H70 did about 20 percent better on SPECweb96 than Sun's Enterprise 450, but trailed the HP N4000 and an

Intel-based Dell server running Windows 2000 by about the same margin. Thus, the CPU performance of the system would be viewed as average. Considering the system's strong HA features, however, we rate the overall performance of the HA-H70 as good.

The HA-H70 is an able candidate for applications requiring server-level HA features, even though HA should not be considered the panacea for all situations. For example, for a Web site with relatively small data requirements, where the data can be easily replicated across a bank of smaller Web servers, a load-balancing appliance that maintains user connections might be a more effective solution than HA. In contrast, the structured and configurable failover features of HACMP fit well in structured application environments such as transaction-oriented e-commerce, where a sophisticated DBMS forms the system's back end. For those more complex environments, the HA-H70 is an appropriate choice, and earns an overall rating of good—three flags.

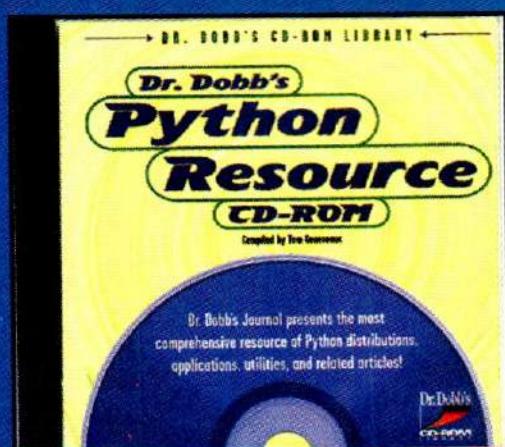
Special Discount For Performance Computing Readers! Save \$10 off all DDJ CD-ROM products!

Dr. Dobb's Python Resource CD-ROM

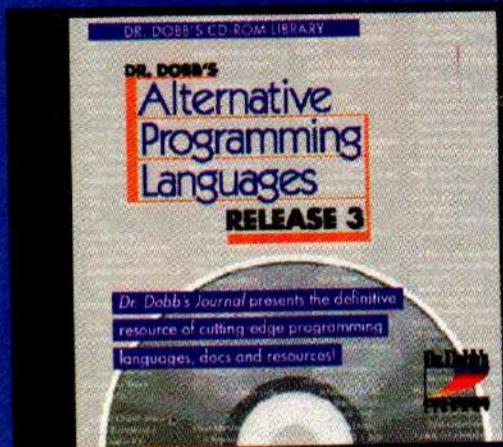
Dr. Dobb's Python Resource CD-ROM is the best available source of Python information you'll find! It includes Python distributions for platforms ranging from Linux to Win32, and utilities and applications for programmers and end users alike. The CD-ROM also includes tutorials, FAQs, and Python articles from leading magazines such as *Dr. Dobb's Journal*, *Web Techniques*, *Linux Journal*, and more!

Python is a powerful interpreted, object-oriented, and portable scripting language that is easily integrated to all of your programming projects.

HTML!



NEW RELEASE! PRICE! \$39.95



Dr. Dobb's Alternative Programming Languages Release 3 CD-ROM

Dr. Dobb's Alternative Programming Languages CD-ROM is back by popular demand! The *Alternative Programming Languages Release 3 CD-ROM* is your definitive resource for cutting-edge programming environments, giving you access to the most innovative, solution-providing languages and tools available today.

This new release contains the most comprehensive collection of programming languages ever assembled with over 600 megabytes of data!

The *Alternative Programming Languages Release 3 CD-ROM* will get you up to speed with languages and docs for Squeak, Dylan, Mozart, Octave, SatherK, Haskell, Lua, EiC, Small Eiffel, Erlang, Modula3, Standard ML of New Jersey and more!

HTML!

NEW RELEASE! PRICE! \$49.95

Dr. Dobb's Tcl/Tk Resource CD-ROM

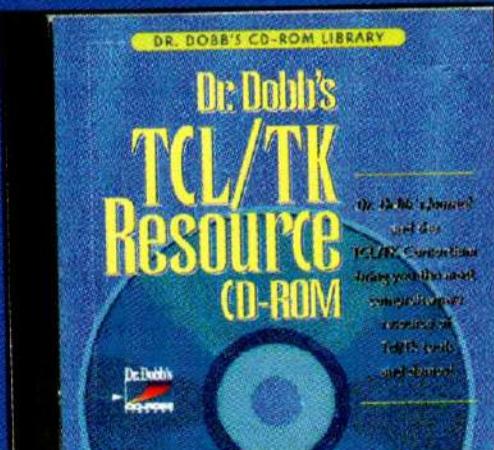
Dr. Dobb's Journal & the *Tcl/Tk Consortium* deliver the most comprehensive resource of *Tcl/Tk* available on one *CD-ROM*! *Tcl/Tk* and its extensions are powerful tools for software developers.

Dr. Dobb's Tcl/Tk Resource CD-ROM is perfect for all programmers who need to get their projects out the door - fast! This *CD-ROM* contains programs and libraries that are ready to run, and easy to install.

PLATFORMS INCLUDED: AIX / Alpha OSF / FreeBSD / HPUX / IRIX / Linux / Macintosh / SCO Solaris / Windows 95/NT /... and more!

USE *TCL/TK* TO:

Build web sites • Perform regression testing • Act as an embedded command language • Create new applications by integrating and extending existing components and applications • Shrink development time with this compact and efficient programming language

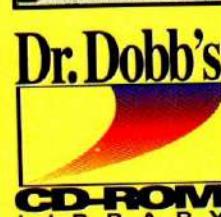


PRICE! \$24.95

HTML!

Visit our web site for all the products from *Dr. Dobb's CD-ROM Library*:

www.ddj.com/cdrom/value/



ORDER TODAY!
800-992-0549
(U.S. & Canada)

All Other Countries: 785-841-1631

• E-mail: orders@mfi.com

• Fax Orders: 785-841-2624

Mail Orders: Dr. Dobb's CD-ROM Library, 1601 West 23rd St., Suite 200, Lawrence, KS 66046-2703 USA

Emmett Dulaney

RED HAT LINUX 6.0

New features add business appeal, but enough for corporate clients?

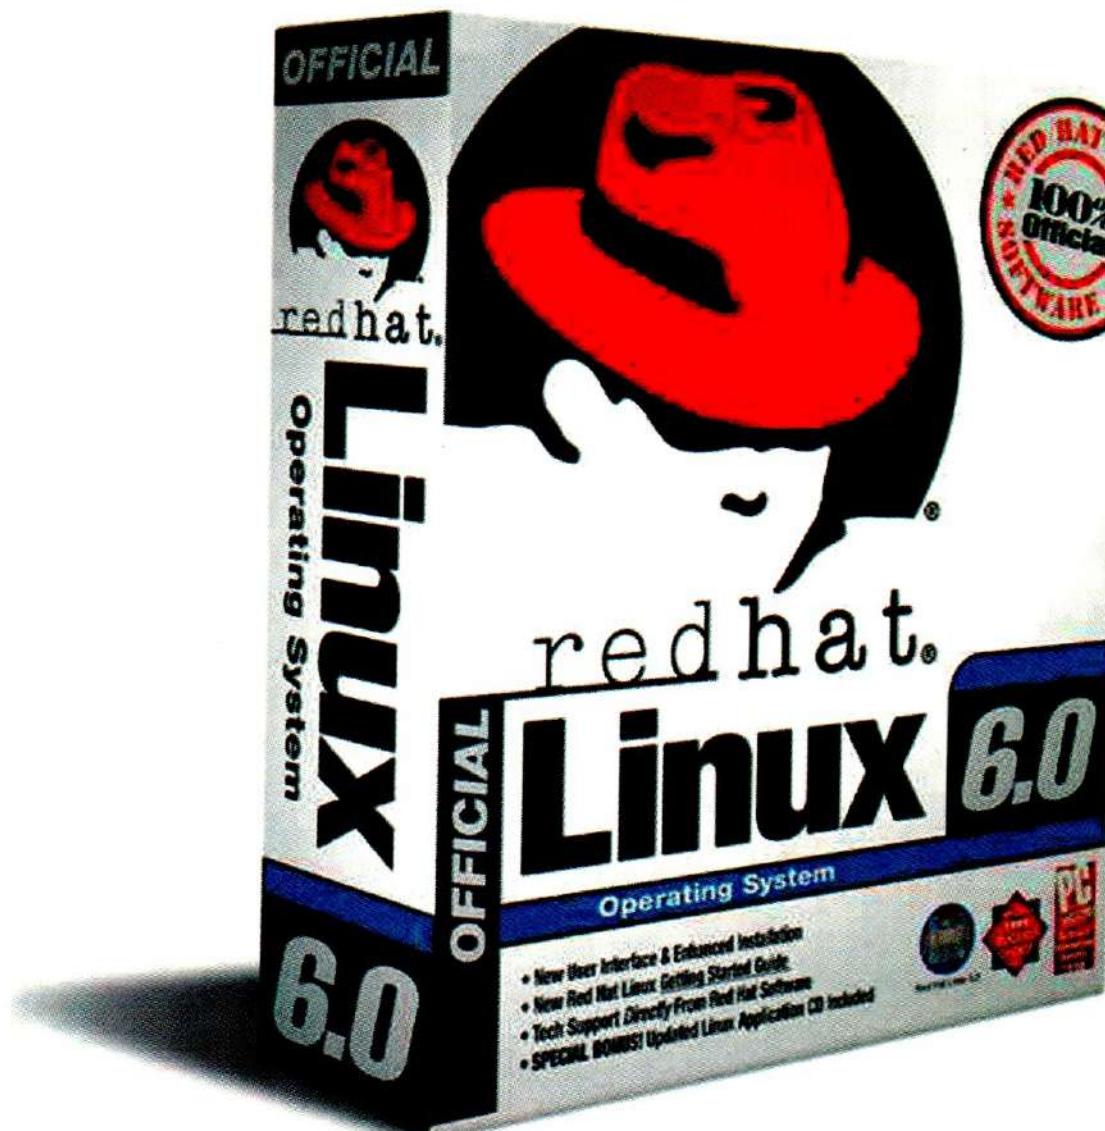
While there are certainly extreme views at both ends of the Linux spectrum, the middle, commercial application of the operating system, is the most interesting. In this arena, Red Hat's Linux distribution is one of the leaders, and garners considerable attention within business and financial operations. In fact, Red Hat Linux has been awarded *InfoWorld's* Operating System of the Year for three consecutive years. Our question, of course, was whether the new release, Red Hat Linux 6.0, could live up to its award-winning heritage.

WHAT MAKES IT SPECIAL

One of the biggest bullets fired from the anti-Linux crowd is that it is an OS without applications. While an OS can be a novelty, the inability to do anything with it can be quite a hindrance. Red Hat addresses this issue head-on by including a CD containing almost 400 applications. Among the applications included are Applixware (an office suite for documents, spreadsheets, and HTML authoring), Flashpoint (C++ to HTML

generator), Star Office (spreadsheets, word processing, e-mail, HTML, and so on), ViaVoice (from IBM, for voice dictation) and VSI-Fax (faxing software). Some of the applications are full versions, others are limited ver-

sions, and still others are simply demos. It is beneficial to have these applications available for an immediately productive use of Linux. At the same time, none of the included applications are compelling enough to



warrant an office-wide conversion to Linux.

Further differentiating itself from the Linux populace, Red Hat provides easy installation, and includes a 300-page Getting Started guide along with a 400-page Installation Guide. Full source code is included on yet another CD (bringing the total to three).

Red Hat also stands out from other distributions of Linux by offering 30 days of full telephone support along with the printed manual and Getting Started guide. In addition, there is 90-day e-mail and fax support.

FEATURE SET

Two user interfaces (GNOME and KDE) provide a Windows-like feel, complete with drag-and-drop capabilities, icons, shortcuts, and a toolbar. If you have ever used Windows 95, learning GNOME is easy.

The Red Hat Package Manager (RPM) 3.0 keeps track of all installed packages: updating, upgrading, and recompiling them as needed. RAID (Redundant Array of Inexpensive Disks) is supported (including hot-swappable features) at the following RAID levels: 0 (no fault tolerance, but the ability to use volume and stripe sets), 1 (fault tolerance as disk mirroring), 4 (dedicated parity drive), and 5 (disk striping with parity).

The Experimental GNU Compiler System (EGCS) is included to simplify development. A vast improvement over earlier C++ compilers, this one supports 64-bit code, namespaces, thread-safe exception handling, and a host of other features.

Also included in Red Hat are Anon-ftp (for anonymous ftp access), Bash (the GNU Borne Again shell) Dhcp-2 (for Dynamic Host [TCP/IP] configuration), ElectricFence (memory allocation debugger), ImageMagick (for displaying X images), Kaffe (a Java virtual machine), Mgetty (for data and fax modems), and Xfree86 (an X server).



Red Hat Linux 6.0's desktop

INSTALLATION

I tested the OS by installing it on several machines. In all cases, I found it to be one of the simplest and most straightforward installs of any OS. The first attempt took 51 minutes, only because I purposely made wrong choices, and rebooted a number of times. Afterwards (knowing what to expect) it was possible to repeatedly complete installations in under 30 minutes.

During the installation, you choose whether you want to install/optimize the OS for use as a workstation, server, or something in-between (custom). The shipping box lists the hardware requirements as 16MB RAM and 500MB free hard-drive space, but you should have (at a minimum) twice that much RAM to even consider the install. The free hard-drive space requirement differs on the type of install you choose: 250MB is the bare-bones minimum for use as a workstation, 600MB is needed for an Internet-ready workstation, 350MB is the bare-bones minimum for the server configuration, and 1.6GB is the recommended minimum for a true general-purpose server.

During the installation, you create an emergency boot disk. A word to

the wise, be careful not to lose this, as it is needed to boot the system if start up files become corrupted. I highly recommend making a copy and storing it in a safe location.

USAGE OF THE 2.2 KERNEL

Many of the features v. 6.0 offers over previous Red Hat versions come as a result of the upgrade to the 2.2 kernel. It is this kernel that provides such features as CODA (file-system support), ipchains (firewall tool replacement for ipfwadm), IPv6 support (also known as IP Next Generation long-file name support), Memory support of up to 1GB on Intel platform automatically detected (2GB through some configuration), Serial consoles, and Swap drive support up to 2GB.

Linux has new versions available almost daily. As I write this, v. 2.2 is considered to be the latest "stable" release, but the versions are now up to 2.3.13. Most of the changes have been minor bug fixes to individual utilities and features. To see what each version's changes are, check the Web at <http://www.kernelnotes.org>.

The next major kernel revision is expected to be 2.4, with a development-completion date expected at the end of the year. Expected features of

2.4 include PCMCIA support, and plug-and-play capabilities.

CERTIFICATION AS WELL

One of the differentiating features between Red Hat and most other Linux vendors is it offers both a training and certification program.

Certification has become the hot commodity in the administrative world, thanks in large part to the MCSE program from Microsoft and the CNE program from Novell. Red Hat's training is different from, say, Microsoft and Novell, because they rely upon multiple-choice exams to certify an engineer/administrator, while Red Hat requires passing a supervised hands-on lab.

To become a Red Hat Certified Engineer (RHCE), one must take prerequisite courses through Red Hat or Global Knowledge Network (<http://www.am.globalknowledge.com>), or have the experience necessary to meet the objectives for the following three courses: Introduction to Red Hat Linux, Red Hat Linux System Administration, and Red Hat Linux Networking and Security Administration.

Each class is four days in length and most include one hands-on workstation per student (in some cases the introduction and system-administration classes can be broken into two segments, I and II, each lasting two days.) Classes are held in North Carolina by Red Hat, and around the country by Global Knowledge.

The price through Red Hat is \$1,998, \$2,098, and \$2,198 for the three classes respectively, and slightly higher at Global Knowledge. From these classes, you are expected to master basic network security, basic security, basic TCP/IP networking, Intel architecture, networking concepts and technology, networking services, UNIX file systems, and UNIX system administration.

With this knowledge, the administrator can then take the RHCE Training and Certification course or choose

to take the exam. RH300 is a five-day course that includes a hands-on exam. If you already know the requirements and do not want to take the class, you can opt for the one-day hands-on lab.

During the actual exam/lab, you are expected to set up and administer a server. You are measured by certain criteria, including security and network services. A proficiency in troubleshooting and diagnostics is also required, as well as a mastery of some common administrative tools.

THE PRODUCT

Red Hat v. 6.0 retails for \$79.95 for the Intel version or \$99.95 if bundled with Power Tools. Red Hat is available as the default OS on new systems purchased from several major vendors including Compaq, Dell, HP, and IBM. It offers full POSIX compliance, symmetric multiprocessing (up to four processors), the NFS file system, and is optimized for Web, intranet, and Internet use. Version 6.0 is based on the 2.2 Linux kernel, supports 2GB swap files, has serial console support, a scalable kernel, and frame buffer support native to the kernel.

As of print time Red Hat just released the 6.0 E-Commerce Server. Security enhancements and features have been added to the base product to provide safer Web-based transactions. Red Hat also markets two other products: Extreme Linux and Maximum RPM.

SHOULD YOU USE IT

If you are looking to install an OS on the laptops of your mobile sales force, or to place on office worker's machines you might want to hold off. This is not because of anything related to Red Hat, but due to the lack of applications for these types of operations. This is almost guaranteed to change in coming months, with Linux versions of Lotus's Domino and other major programs in the works.

If you want to stay current try the "RHMember Program," a subscription service through Red Hat. Members are automatically shipped every new update for a year, ensuring you always have the latest version of Linux. The enhanced "RHMember More Program" works on the same concept, where instead of being sent updates, they are compiled on CDs and mailed every six weeks.

If you are looking for a robust OS that can be used for Web operations, then Red Hat Linux is an excellent choice. Red Hat is also an excellent choice if you are looking for an OS to be used by your developers to replace much more expensive UNIX machines, or any similar operation.

Contact Red Hat Software Inc., P.O. Box 13588 Research Triangle Park, NC 27709. Tel.: 919-547-0012 or 888-RED-HAT1, fax: 919-547-0024, Web: <http://www.redhat.com>.

Red Hat Linux 6.0 Web resources

To learn more about Red Hat Linux, or to purchase for installation on your systems, check the following sites:

- Development
- General Linux Information
- Hardware compatibility
- Home Page
- Knowledgebase
- Purchasing
- Support
- Training/Certification

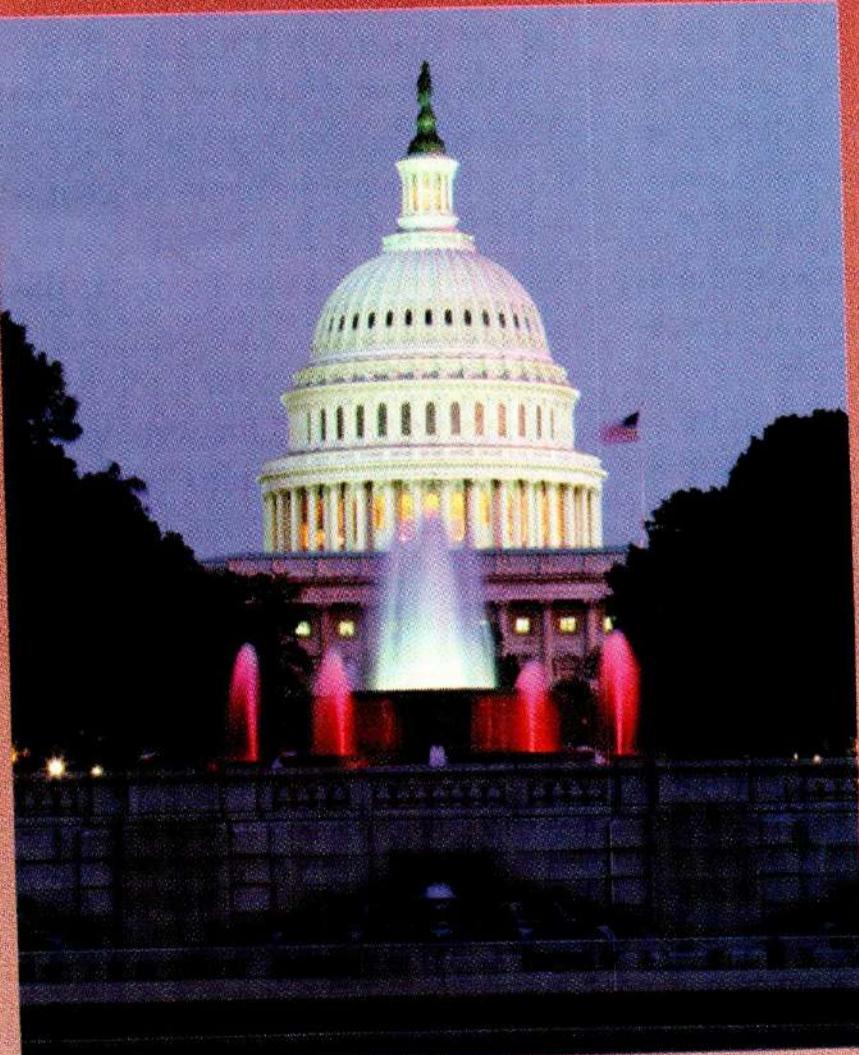
- <http://developer.redhat.com>
- <http://www.linux.com> or <http://www.linux.org>
- <http://support.redhat.com/hardware/>
- <http://www.redhat.com>
- <http://www.redhat.com/knowledgebase/kb/cache/1.html> or
- <http://www.redhat.com/corp/support/errata/index.html>
- <http://store.redhat.com/commerce/>
- <http://support.redhat.com>
- <http://www.redhat.com/products/training.html>

26th Annual Computer Security Conference & Exhibition

November 15-17, 1999

Washington, D.C.

Marriot Wardman Park



Choose from over 130 sessions
and topics that include:

Network Attacks and
Countermeasures
Windows NT Security
Secure Web Commerce
Intrusion Detection and
Response
Cryptography
Hackers
Security Awareness
VPNs and Firewalls
Selling Security to Management
Policies and Procedures
—and much more

“THE Security Event of the Year”

FAX THIS BACK TO CSI AT: 415-905-2218. FOR MORE INFORMATION CALL: 415-905-2626. EMAIL: csi@mfi.com

- YES! Please send information on attending.
 YES! Please send information on exhibiting.

Name

Title

Company

Address

City

State

Zip

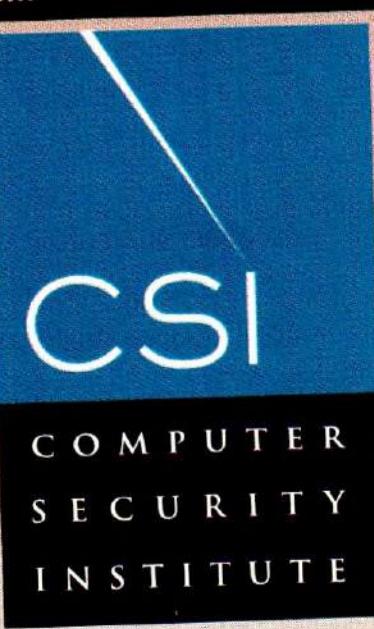
Country

Phone

Fax

E-mail

PFC 10/99



www.gocsi.com

Ben Thompson

ACCRUE INSIGHT 3.0

Accrue Insight 3.0 provides flexibility in measuring Web site traffic and analyzing visit patterns.

In the Web world, a site's value and popularity can be distilled down into numbers—how many hits, the duration of page views, and how many users clicked through the banner ads. Site administrators need a way to keep an accurate accounting of these numbers and generate informative reports. Accrue Insight, a Web traffic measurement and analysis tool, is one such way.

If you only need basic statistics such as aggregate visits, files served, and page views to your site, then Accrue Insight may be more than you need. This software is aimed at administrators who want to perform detailed performance analyses of their Web sites, rather than network engineers planning network capacity, although Insight can be used for both.

We used Insight to measure the performance of individual sites as well as retrieve aggregate traffic in a Web farm of over 180 domains. Our software requirements included the ability to measure server response time and allow system users on different continents to organize data in myriad ways. We knew we needed a packet sniffer on our network and a

good Web-based user-authenticated reporting tool on the front end for system users. And we knew that robust reporting would require that data be stored in a high-caliber database.

Accrue Software designed Insight with many of these features. Perhaps the most important improvement in version 3.0 (released in April 1999) is the improved database loading speed. After upgrading the warehouse software, nightly load times on our systems were reduced from over six hours to under two (for the over-180 domains mentioned earlier).

IN THE FIELD

Insight is designed for large-scale enterprise Web environments. The bundled Red Brick (now part of Informix) RDBMS performs well under large volumes, but is memory intensive. A Sun E450 with 512MB RAM often paged out (swapping was never observed) and sometimes blocked processor access during simultaneous ad hoc queries against a database containing over 1.3 million HTTP requests per day. However, even when blocking processor access due to

memory limits, the processors barely broke a sweat, with logs showing a peak system load of only 1.5. Batch-report generation is recommended for all repetitive reporting needs and is less memory intensive.

Much of Insight's power, besides the sheer scale of the system, is provided by the Report Wizard (see Figure 1, p. 52), a reporting interface that allows authorized users to sort data in various ways. The Wizard is a later addition to Insight and is slowly replacing the earlier Java-applet-based Report Explorer. Unfortunately, the reporting functionality is currently split between the two interfaces. For example, the Wizard allows users to view referrals from user-specified domains (including query strings), but not referrals to specified URLs, while Explorer provides a means of viewing top referrers to specified URLs, but not from specified referrers. In addition, when using the Wizard, a user can view the total time spent (by all visitors) to a page or a site, and the number of page views and visits, but it doesn't do the simple math to give the average time spent per page or per visit, which is aggravating. The

Explorer does provide a report on average time spent per page, but not per visit. Earlier annoyances in the 2.6 release were resolved in the current version, and we have been solicited by Accrue for product suggestions twice in six months.

Insight ships with an extensive set of predefined reports in both the Wizard and Explorer (see Figure 2). The Explorer reports can be set up

through filters made up of groups of site URLs and/or client hosts. Using URL group filters, a Web publisher can run reports on specific areas of a site. With host filters, site publishers can see, for example, which pages of their site their competitors or their customers are viewing. However, this flexibility comes at a price. When filters are used, the software must read

though few installations are tracking 180 domains, our total number of requests per day is not atypical.

Insight does not ship with extensive charting or graphing capabilities, but instead allows all reports, both from Wizard and Explorer, to be output as comma-delimited text for easy import to a spreadsheet application. This is a strategic decision by Accrue, and it makes sense, since engineering talent is better spent on core data-crunching issues, rather than reinventing graphics capabilities every enterprise already has.

A major boon to Explorer is a navigation graph that provides a flow chart of top paths to and through a site (see Figure 3). Users may designate any page of their site as the "focus node" and view referrals to and redirects from that node. A simple click on a different node makes that node the focus. Host and URL group filters may be applied to the navigation graph as well.

The network traffic-collector piece can be run on the same box as the data warehouse or separated. An advantage to running them apart is that the collector can hold raw data for several days (or weeks if space is available) in case the warehouse goes offline. The collector writes individual HTTP requests into domain-specific directories, creating a new file every hour. Typically the Insight warehouse pulls data from the collector every hour. It checks to see which of the files on the collector have already been moved over, then moves the new ones. If you are saving data on the collector for several days, and have many domains, the warehouse must check thousands of files just to move over a handful. When we held data on the collector for eight days, we found that transfer times varied from 35 to 60 minutes for one hour of data, allowing no breathing room for increases in traffic. Cutting the time we held data on the collector resulted in a linear decrease in transfer time. Most enterprises don't run scores of sites, so this should not be a problem.

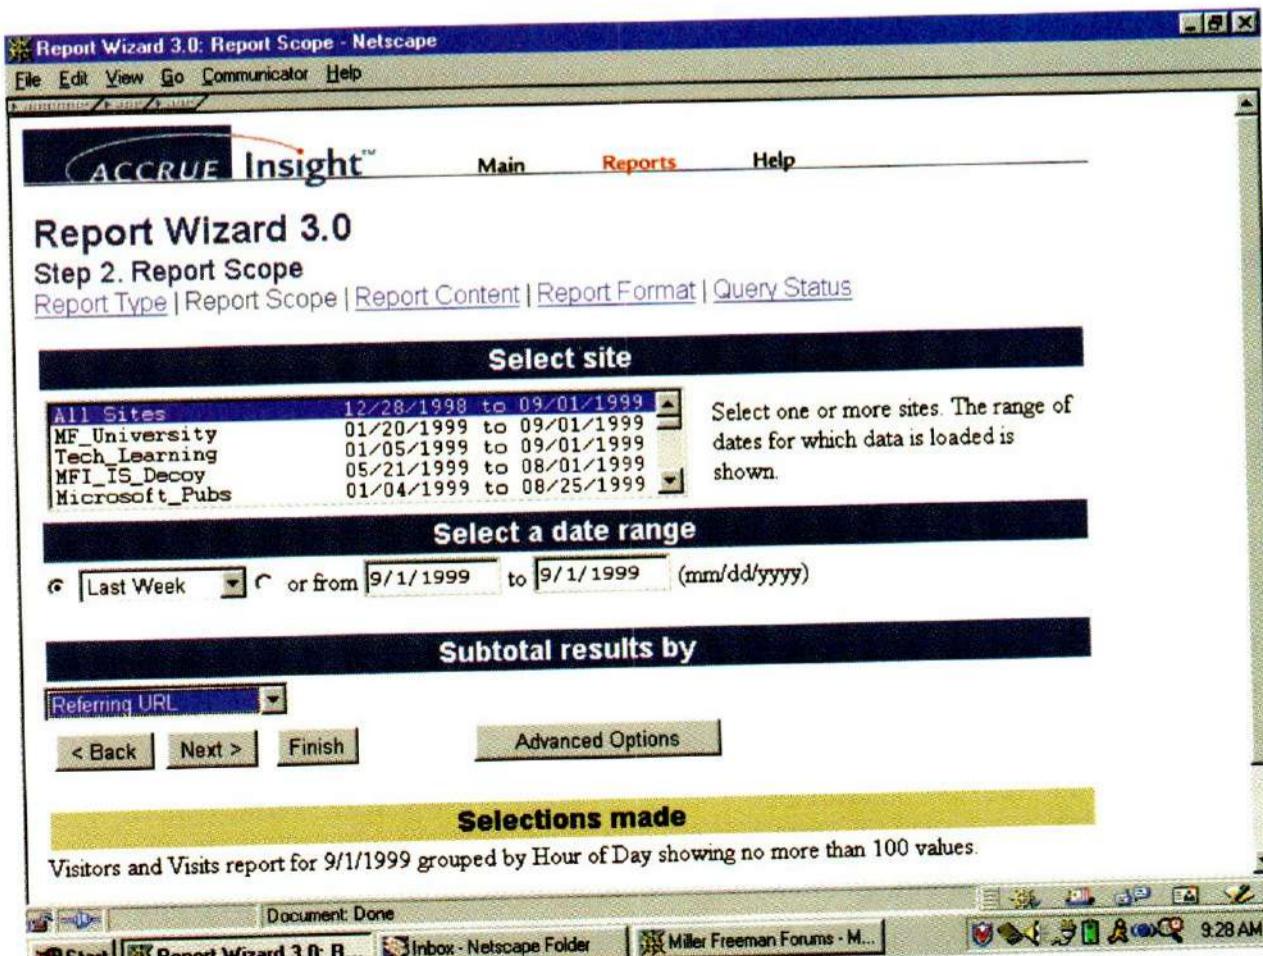
The package is highly configurable, letting application administrators specify MIME types to capture.

to run in overnight batch mode and be e-mailed to marketers and/or posted to personal report pages. The reports available through the Wizard cannot currently be batched, but Accrue is planning on the next major product revision to include this feature.

Figure 1 The Insight Report Wizard allows numerous user options for generating detailed reports. Report configurations can be saved by users for later re-use.

through the detailed hit table, which can boost the time it takes to get report results from seconds to minutes.

Generating queries that cover a long time span also boosts the wait time, even when not using filters. When we tried to run reports showing top-level hosts making requests to our combined domains, we could not specify time spans longer than ten days without hanging the system. Al-



INSTALLATION

Installation and setup of the software is straightforward but should be done with an administrator's support. The package is highly configurable, letting application administrators specify MIME types to capture (that is, true pages, not sub-elements within the page), hosts to ignore (for example, internal traffic), and whether to grab query strings (and from which referrers). Also, data from Web server access files can be auto-converted to Accrue format and uploaded to the warehouse. All that is required to write the necessary configuration files is a basic understanding of regular expressions. Again, however, there is a cost associated with gathering the data that is no reflection on the product: capturing query string data fills up the referrer table more quickly, since there are potentially many more unique referrers. Conversely, excluding image requests from the detailed hit table saves not only disk space, but (more importantly) reduces query times.

The table schema is arranged to store requests for content that may be split among (or mirrored across) several servers. Each requested resource

Figure 2 "MyAccrue" provides individual end users with a set of reports that are batch generated nightly. Reports can also be e-mailed automatically. A default report set or custom reports can be run for each user.

is assigned to a specified "content set" (typically a single domain). Each content set can be associated with one or several servers, or one server can be associated with several content sets. Likewise, a site can comprise one or more content sets.

Several data-management tools are provided through the admin user interface, including data-segment cycling. This lets the administrator toss out older detailed data (with the option to retain summary data), recycle disk space, and delete data for selected sites entirely. Unfortunately, data for specific content sets cannot be so easily deleted. If your content set/site/server relationships change over time, reflecting those changes in the Insight database can require creative SQL gymnastics.

Insight documentation is extensive but the editing is poor. Some features are not documented at all, such as the command-line syntax for performing

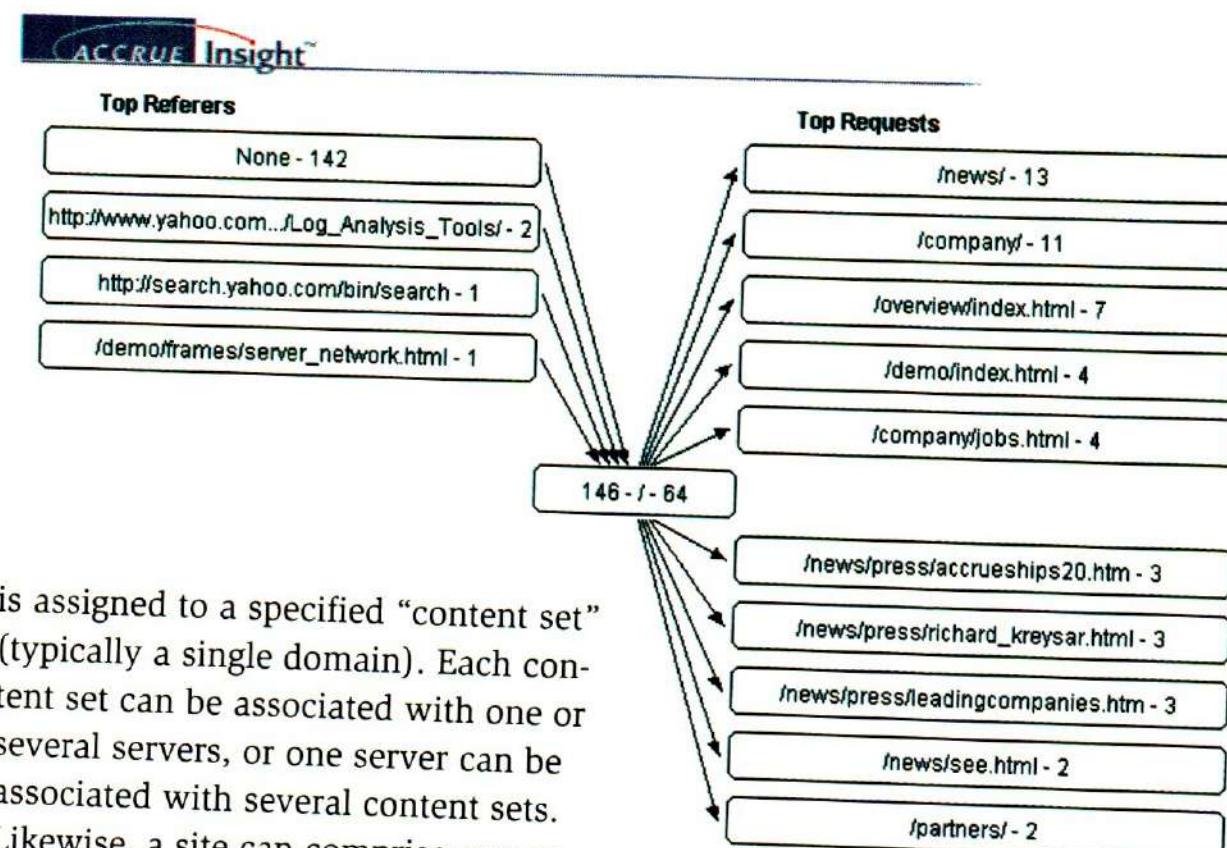


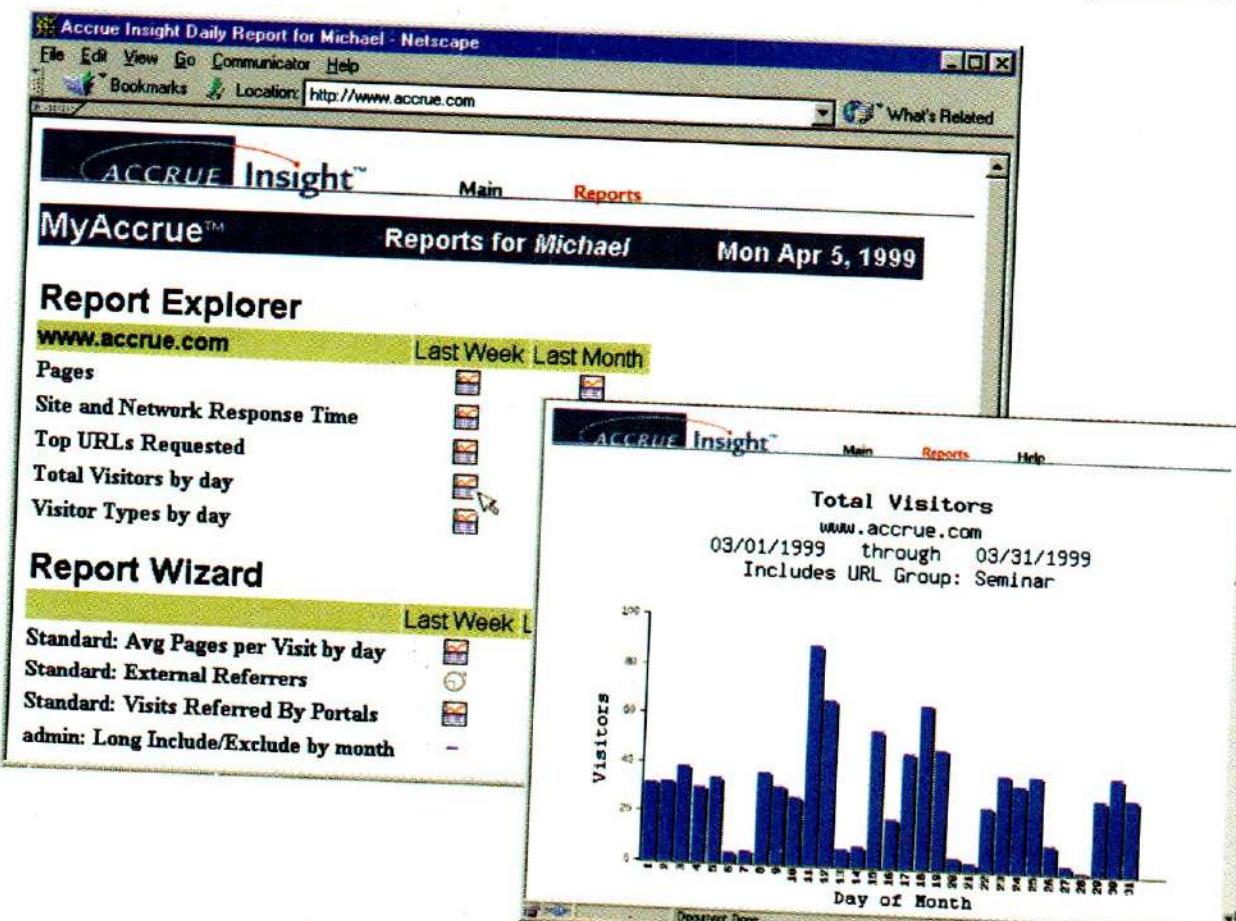
Figure 3 The NavGraph shows the path users have taken through a client's Web site.

manual data loads. Tech support is in line with other small software firms undergoing growing pains. The technicians are competent and know their way around the product, but can be stumped occasionally. The principle drawback is lack of an escalation path. If the front line support can't answer, then you either pay for an engineer to make a site visit or you figure things out yourself.

There are at least a half-dozen other software package vendors in the Web site traffic-analysis space, with two or three others in the enterprise class. Accrue Insight should make any short list, though, due to its basic architecture, scalability, and flexibility.

Contact Accrue Software Inc., 48634 Milmont Dr., Fremont, CA 94538. Tel.: 510-580-4500 or 888-4ACCRUE, fax: 510-580-4501, Web: <http://www.accrue.com>.

Supported platforms: Datastore (warehouse) requires Solaris 2.5/2.6, Network Collectors require Solaris 2.5/2.6 or Windows NT 4.0. A Java-enabled Web browser is required for using the UI (Netscape Navigator 3.0.1-4.5 or Microsoft Internet Explorer 4.x recommended). Note: Accrue is reportedly close to releasing a version of their warehouse ported to Oracle 8.x.





BRIDGING THE [R]EVOLUTION OF THE LINUX SOLUTION

Jacob Javits Center

New York City

Conference: February 1- 4, 2000

Expo: February 2 - 4, 2000

Take control of your operating system, the power of innovation is in your hands! Linux users, software developers, and business leaders, participate in the evolution of the enterprise.

LinuxWorld Conference & Expo in New York City is the premiere event, for supporters, investors and those interested in the linux operating system. Hosted in the financial capital of the world, the knowledge shared at this event will reinvent the way technology and business function in the new millennium.

For more information on exhibiting,
please contact:
LinuxWorld Conference & Expo
3 Speen Street, Suite 320
Framingham, MA 01701

For more information or to register,
please visit:
www.linuxworldexpo.com
or call 800.647.1474

Owned and Managed By



www.linuxworldexpo.com

Fla
Sp

Lin
W

Media Sp

LIN
THE CHRONIC

UNIX
PERFORMA
COMPUT

Sy

INFO
WORLD Tech Smart
Business S

Networ

Dr. Dobb's
JOURNAL SOFTWARE
TOOLS FOR THE
PROFESSIONAL
PROGRAMMER

Cornerstone S

VA LINUX
Systems

Platinum S

LINUXCAR

SYBASE
INFORMATION ANYWHERE.

IGEL



William LeFebvre

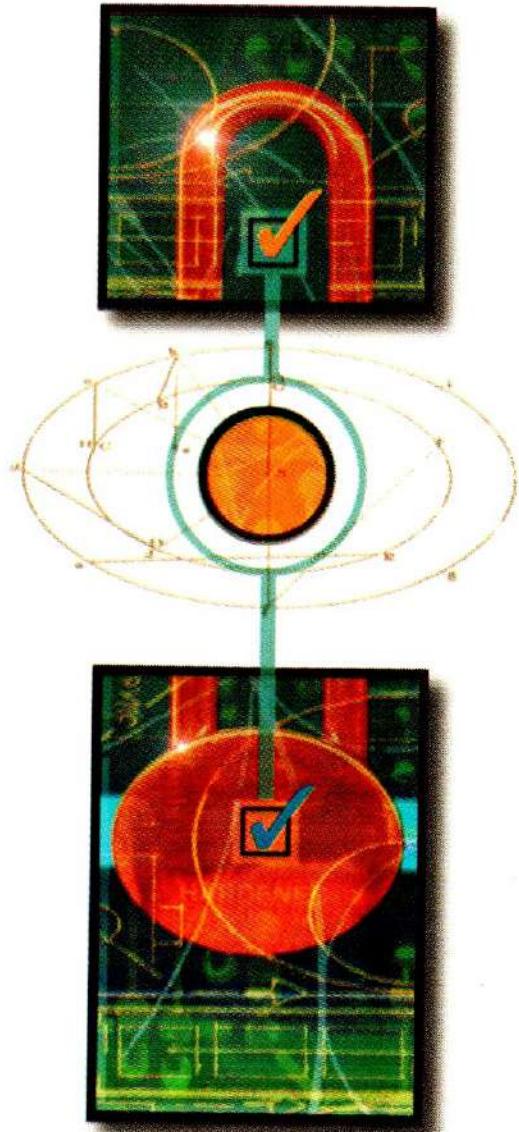
INTERNET BLACK HOLES

When the founding fathers of the Internet created a mechanism for the exchange of electronic mail, little did they realize that they were inventing a new marketing channel. Network old-timers like myself have always felt it is rude and inappropriate to send unsolicited and unwanted electronic mail to someone. When I first started receiving these "spam" messages I was shocked and a little angered. My immediate response was to reply and ask the individual why he sent that particular message to me. Of course, I didn't receive a reply.

Spam has been a problem on the Internet for most of the '90s. The feeling among administrators has always been one of near hopelessness. How does one differentiate between legitimate mail and unwanted junk? Recent releases of Sendmail have included mechanisms for detecting badly formed mail messages, a hallmark of spammers (see "Sendmail And Spam," August 1998, p. 55). One ability Sendmail now boasts is a blacklist: a file listing sites from which you never want to accept mail.

Blacklisting is certainly one way of limiting spam—any site that has a history of originating spam can be added to the list. But there are many sites, and the possibility of a spam-hosting site repenting of its ways. Who is going to maintain the blacklist?

The Mail Abuse Prevention System, LLC (MAPS), a California non-profit corporation, has an answer. They call it the Realtime Blackhole List (RBL). This is a list of sites known to originate or relay spam. Individuals at MAPS maintain this list for the benefit of their mail servers and for anyone else who wants to use the information.



MAPS' VIEW OF SPAM

The individuals who formed MAPS are fed up with spam—unsolicited messages are unwelcome and they clog networks, servers, disks, and mailboxes, which led MAPS to conclude spam is nothing short of theft of service. Those who send the messages have found ways to minimize or eliminate their network-access cost, and the unwilling recipients bear the brunt of the costs of transporting and storing the spam.

But what exactly qualifies as spam? In the old days, electronic mail was a personal medium. Messages you received were usually from people you knew or had a passing acquaintance with. Those who weren't were asking you for specific information or sending something they knew would interest you. In other words, the mail was of direct benefit to both sender and recipient.

Spam is anything we don't want, that was not targeted specifically for us. It is electronic junk mail, usually characterized by mass mailings that use bogus return addresses and methods of delivery that are designed to minimize costs to the sender. I have worked with people that maintain mailing lists of customers, and have heard the term "spamming" used to describe the act of mass mailing to such a list. I don't agree with that usage. Presumably, the members of this list requested to be on it, or at the least did not request to be left off.

The members of the list have an easy mechanism for removal, and requests to be removed are honored. Finally, those sending out the mass mailings place the burden of delivery on their own equipment, rather than dumping it on some unsuspecting relay host.

It turns out there are two ways to facilitate the delivery of spam. One is to play host to spammers by providing them with the facilities to inject their messages into the Internet. Usually this is an Internet Service Provider

(ISP) who does not know or does not care that one of its customers is sending out massive quantities of unsolicited mail. The other way to help spam along is a passive roll—letting your mail servers act as a relay point for messages. This can happen to anyone who runs a mail server directly connected to the Internet. If you are not running a mail transport agent (MTA) that specifically blocks relaying, then your server could be used by spammers as a place to fan out their messages: they send you one message with a thousand recipients, and your server does all the work.

THE REALTIME BLACKHOLE LIST

MAPS maintains a list of sites that, in its view, participate in propagating spam. The primary purpose of this list is to protect MAPS' own servers from spam attacks. Since it views spam as theft of service, it feels it is within its rights to protect access to their service. This happens through the Realtime Blackhole List (RBL). Since accumulating and tracking sites is a monumental task, other sites became interested in gaining access to the RBL so they wouldn't have to duplicate the effort.

It is important to note that use of the RBL is completely voluntary. As a site administrator or network service provider, you can choose to use the RBL for filtering or you can ignore it. This has kept MAPS out of legal trouble from corporations that found themselves added to the list. There is a certain amount of trust involved in utilizing the RBL: you are allowing an outside organization to control filtering of your incoming mail or IP packets.

HOW THE RBL WORKS

The RBL is a list of IP addresses. It does not contain domain names or mail addresses, just IP addresses. If MAPS discovers that there is a lot of spam originating from a specific address, or within a block of addresses, the addresses will get added to the list. Anyone utilizing the list will automatically stop accepting mail—any mail from the addresses. Some sites have gone a step further and have configured their routers to drop any packets coming from those addresses, whether they are SMTP packets or not.

There are two separate manifestations of the RBL. The first and most accessible is glued onto the Domain Name System (DNS). The domain *tbl.maps.vix.com* has been set aside to contain IP addresses currently on the list. These addresses are indexed like they are in the reverse domain space (*in-addr.arpa*), with the four octets reversed. Therefore, if 10.1.2.3 was in the list, then an address record (A record) would exist for the domain name 3.2.1.10.map.vix.com. All address records in the

RBL refer to the same address: 127.0.0.2. This is just an arbitrary choice, but MAPS says you can count on it remaining this way. The actual address used in the A record is not important: it is the presence of the A record that indicates membership in the RBL. For testing purposes, the address 127.0.0.2 is in the RBL and 127.0.0.1 is not. Try a DNS lookup for *2.0.0.127.tbl.maps.vix.com* and see what you get. In this form, the list can be consulted by any mail transfer agent (such as Sendmail) through a DNS lookup. If the A record exists, then the mail is bounced.

The second manifestation of the RBL is in the form of Border Gateway Protocol (BGP) routing advertisements. The BGP advertisements are there for service providers and others that run their own "autonomous systems." They can utilize the information to route all traffic from any listed network address into the bit bucket. This is truly a black hole in the fabric of the Internet.

Information in the RBL domain can be accessed in one of two ways. Anyone can get the data through a DNS lookup for a specific address. In this form, the information is freely available to anyone who wishes to use it. The second form of access is only available by subscription from MAPS, and that is in the form of a DNS zone transfer from one of their name servers. The zone transfer is ideal for larger organizations that handle a higher volume of mail. With a local copy of the zone they don't need to do an off-site DNS lookup for every incoming message.

GETTING ON THE RBL

Not that you would actively seek to be placed on the RBL, but you may wonder what sorts of activities would justify an addition to the list. MAPS feels that any site that acts in a way that is spam-friendly deserves to be added. It also feels that sites acting as relays, even inadvertently, need to be added to the RBL until they update their software to block relayed messages. The justification for this strong position is that only by blocking the relay sites can the RBL hope to stop spam. Those who run the list do not add sites easily or quickly. They consider the list a last resort, and will only use it when more diplomatic channels have failed. Initially, efforts are made to contact the responsible parties and service providers to inform them of the problem.

Once an address is added to the RBL, any site that uses it to block mail will stop accepting *any* mail from the address. Some sites are configured to block all traffic from blacklisted addresses. The result is a noticeable decrease in connectivity for the site involved. Where the issue gets sticky is when MAPS resorts to adding an entire block of addresses allocated to a spam-friendly ISP. The connectivity blackout then extends beyond spammers, and ends up affecting innocent third-parties—customers

of the same ISP. But remember, those who are blocking the packets are doing so voluntarily. They have decided to trust the judgment of the folks at MAPS to block sources of spam.

GETTING OFF THE RBL

Getting off the list can be easier than getting on. Once MAPS determines that a listed site has taken steps to change its spam policy, MAPS will remove the site from the list. For ISPs this could mean a change in their access-use policy. A good policy is a service provider's best defense against customers who abuse their Internet access. Many large ISPs now have a strong use policy that prohibits activities like spamming. MAPS is willing to work with ISPs to help them revise and strengthen their use policies towards this end.

Relay sites are removed from the list once MAPS can determine the site no longer permits the relaying of mail messages. The people at MAPS can provide an administrator with resources for strengthening MTA software to guard against undesired relaying.

OPPONENTS

The RBL project is not without its detractors, however. Some have compared it to McCarthyism, and some organizations finding themselves on the list have threatened to sue. To my knowledge, no lawsuit has actually gone to court. The policies used by MAPS are available on its Web site (<http://www.maps.vix.com>) and are quite clear. The site provides a list of activities that will get you added to the RBL, and a list of activities that will get you off. The RBL information is also available via a Web interface, designed for interactive use by you and me. The list is not kept secret: anyone with a Web browser can find out if any specific IP address is on it.

Whether you agree with the efforts of MAPS or not, you need to be aware of it and the RBL. If you want to use the data in the list to further guard against spam, many of the MTAs can be configured to do just that. Agents that have support for accessing the RBL via DNS are NT Mail, Postfix, qmail, Sendmail, smail, and Zmailer to name a few. If you think you've been blacklisted, check the list. If the site indicates your address is on it, the best place to start is with the resources on the Web page, and a phone call to MAPS.

Waiting at the wrong terminal for Web-to-host access?

The Internet and Web-to-host access are changing how business travels. Critical information that was previously accessed by dumb terminals is now available by browser over the Web with TinyTERM products from Century Software.

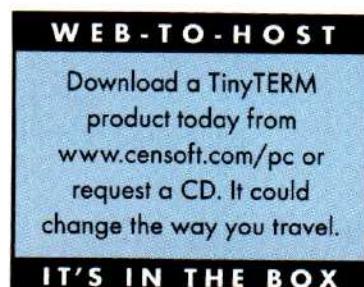
Go Anywhere in the World

Now you can provide access to mission-critical data on your legacy host—inventory reports, for example—over the Internet. Customers can log in, search for what they need, compare prices, locate the nearest available outlet, then order—all on the Web.

Stay Put

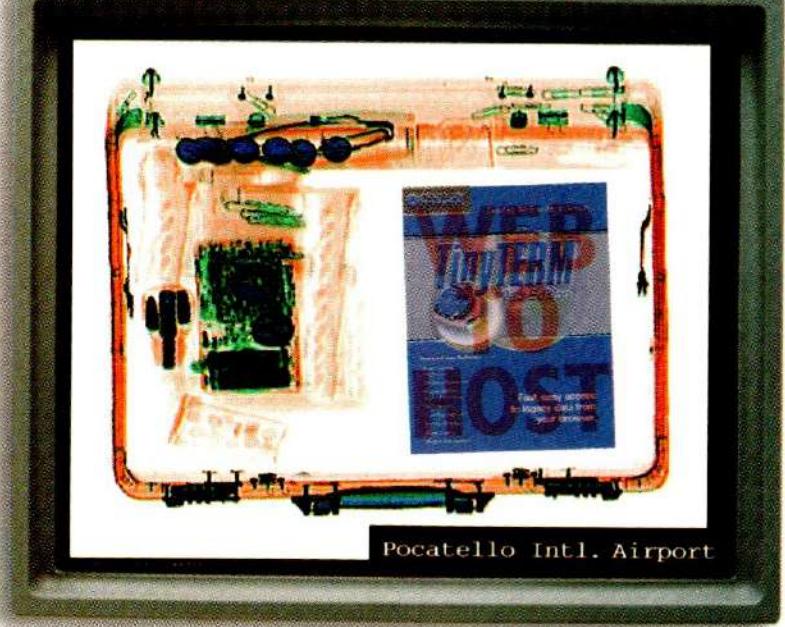
With TinyTERM technology, IS managers don't have to travel to provide global access to your host data. TinyTERM can be installed, configured, and administered from a central location, eliminating the hassle and cost of large-scale deployment.

The power of Web-to-host access is self-service information. It's in the box.



www.centurysoftware.com tel. 800-877-3088 or 801-268-3088 fax 801-268-2772

Select No. 21 at www.performancecomputing.com/fyi/



TinyTERM Web Server Edition

- Provides intranet, extranet, and Internet users access to UNIX data and applications via the Web. Features server-based deployment, installation, and configuration.

TinyTERM Plus Edition

- Provides client access to UNIX and IBM host data and applications through a Web browser interface. Features LPR/LPD and NFS file sharing utilities.

TinyTERM Thin Client Edition

- Provides access to UNIX data and applications in a thin-client/server environment that is compatible with Citrix MetaFrame and Microsoft WTS. Features server-based deployment, installation, and configuration.

CENTURY
SOFTWARE.COM
Leading Web-to-Host Access

Eric Foster-Johnson

HIGH-DEFINITION X COLOR

When the X Window system became popular, particularly on UNIX workstations, you were lucky if you had a display that supported 256 colors. At the time, most X programmers used the simplest color functions because most workstations didn't provide more than the simplest color displays.

In recent years that's all changed and we now expect even low-end hardware to display thousands of colors. Even so, most X programs still stick to old habits and don't take advantage of the greater color capabilities that X provides.

This month we'll go over color programming, starting from setting up colors in a colormap to using more-complicated X visuals.

X builds its color model on RGB components, the red, green, and blue values used to define color on a typical monitor. Using the Xcms routines, short for X color management system, you can define colors in terms of CIE (a color model defined in 1931 that represents all possible colors in a three-dimensional color space)¹ or HSV (hue saturation value) definitions, but the Xcms routines translate either definition to the RGB values the X server requires.

In X, the RGB values go from 0 (all off) to 65535 (all on). Many color definitions that you'll find, though, go from 0 to 255. It's not that hard to convert with a macro such as the following:

```
#define Conv256To64K(v) \
((65535 * (long) v)/256)
```

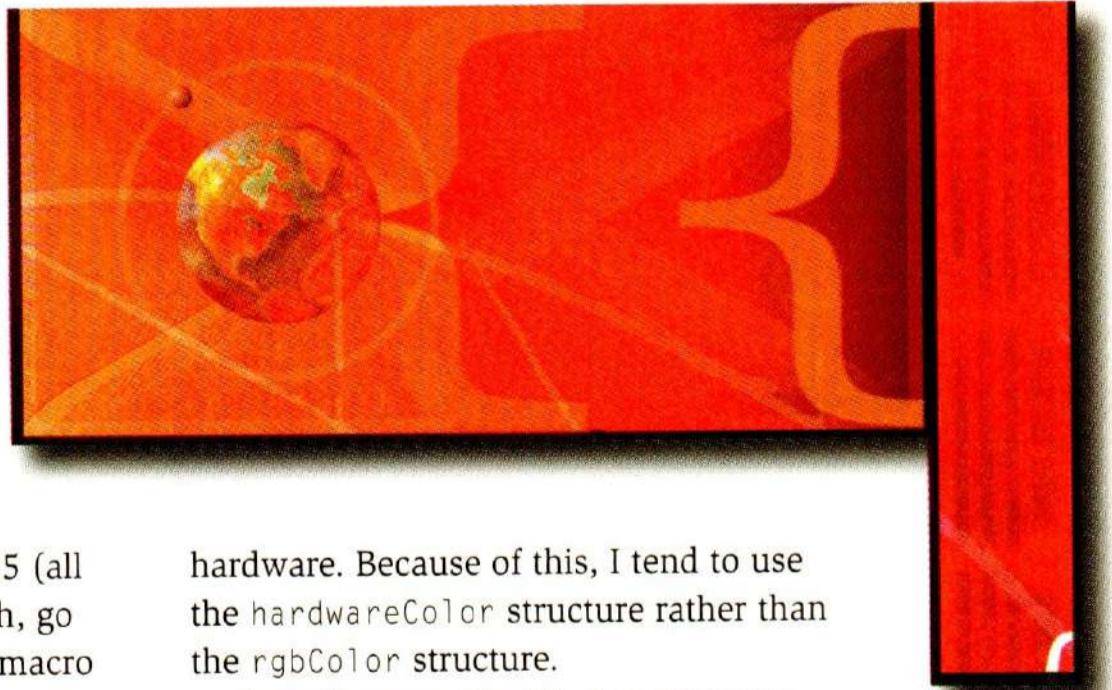
Whatever method you use, if you start with RGB values, you need to follow the X scale of 0 to 65535.

You can allocate color cells from a colormap using RGB values. Or, you can look up colors by name from a color database. X provides a database of more than 700 color names (most of them shades of gray, though). Each color in the database is defined with RGB values (interestingly,

using a range of 0 to 255). You can look up common colors by name with the XLookupColor function:

```
Status XLookupColor(
    Display* display,
    Colormap colormap,
    char* colorName,
    XColor* rgbColor,      /* RETURN */
    XColor* hardwareColor) /* RETURN */
```

On success, XLookupColor returns a nonzero value and fills in the two XColor structures. XLookupColor fills up the rgbColor structure with the red, green, and blue components as read in from the color database. XLookupColor also fills the hardwareColor structure with the closest match to the color definition that is supported on your



hardware. Because of this, I tend to use the hardwareColor structure rather than the rgbColor structure.

In both cases, the XColor structure holds the following values:

```
typedef struct {
    unsigned long pixel;
    unsigned short red, green, blue;
    char flags;
    char pad;
} XColor;
```

XLookupColor fills in the red, green, and blue fields, along with the flags, which indicate which of the fields

were used by bitmasks: DoRed, DoGreen, and DoBlue, respectively.

You can then allocate cells in a colormap from the color definitions held in the XColor structures filled in by XLookupColor. X calls these color cells pixels, which aren't dots on the screen, but values that represent a position in a colormap that holds a given color. In simple colormaps, the pixel is merely an index into an array. This is not always the case though, as we'll discuss later.

To allocate a color cell, call XAllocColor:

```
Status XAllocColor(
    Display* display,
    Colormap colormap,
    XColor* hardwareColor) /* in/out */
```

The XColor pointer is both an input and output parameter. On input, XAllocColor reads the RGB values and flags fields in the structure. On completion, XAllocColor fills in the pixel field with the pixel value for that color in the passed-in colormap.

Before allocating a new color cell, XAllocColor checks if a color cell already holds the given RGB values in the colormap you pass in. This lets applications share common color cells, say, for red or gray. If you pass in the default colormap, chances are most common colors are already allocated.

On success, you can use the pixel field for drawing. With the X library, you can call XSetForeground to store the color into a graphic context as the drawing color:

```
int XSetForeground(
    Display * display,
    GC gc,
    unsigned long pixel)
```

With a Motif application, you can use the pixel value to set the foreground or background resources for a widget. For example:

```
XtVaSetValues(widget,
    XmNforeground,
    hardwareColor.pixel,
    NULL);
```

You can also combine the tasks of looking up a color from a name and allocating a cell with XAllocNamedColor.

Both XAllocColor and XAllocNamedColor allocate read-only color cells. That is, once set, you cannot change the definition of the color in that cell. Because of this, multiple applications, using the same colormap, can share the same colors.

It's slightly more difficult to allocate cells that you can change later, called read-write color cells. To do this, call XAllocColorCells or XAllocColorPlanes. Allocat-

ing color planes lets you do fun things like use color planes for double-buffering. A later column will take that on. For now, XAllocColorCells takes the following parameters:

```
Status XAllocColorCells(
    Display* display,
    Colormap colormap,
    Bool contiguous,
    unsigned long* planeMasks,
    unsigned int numberPlanes,
    unsigned long* pixels,
    unsigned int numberPixels)
```

With XAllocColorCells, you can allocate both color planes and individual cells. If you set the contiguous flag to True, you can insist that the color planes allocated must be contiguous. This is not always possible if a number of colors have already been allocated from the colormap.

If you're just interested in read-write color cells, pass False for the contiguous value, pass NULL for the planeMasks and zero for the numberPlanes. Pass the number of pixels—color cells—you want to allocate in the numberPixels parameter, and an array of unsigned long values long enough to hold the returned pixel IDs in the pixels parameter. For example:

```
unsigned long pixels[12];

status = XAllocColorCells(
    display,
    colormap,
    False,
    (unsigned long*) NULL,
    0,
    pixels,
    12);

if (status != 0) {
    /* We're OK to use the
    colors in the pixels array...*/
}
```

This call defines a 12-element array and asks for 12 color cells.

Allocating read-write color cells is just the first step. X gives you these color cells without defining any colors in the cells. The next step is to do this by calling XStoreColor:

```
XStoreColor(
    Display* display,
    Colormap colormap,
    XColor* colorDefinition)
```

With XStoreColor, you must not only fill in the RGB values in the XColor structure, but the pixel field as well. Set

the pixel field to one of the pixels you got back from the call to `XAllocColorCells`. Also remember to set the flags field, typically to the `DoRed | DoGreen | DoBlue`, the OR of these bitmasks.

Once you store in colors, you can then use these color cells, the pixel values, with normal X drawing or widget color values. You can call `XStoreColor`, or its brethren `XStoreColors` or `XStoreNamedColor`, to then change the definition of a color cell on the fly. This is useful for animated effects or letting a user modify the interface.

COLORMAPS

So far, all of these calls require a colormap. You allocate colors in a colormap. You can create multiple colormaps for different tasks and your system hardware may be able to display one or more colormaps at the same time.

If you use more colormaps than your hardware can support simultaneously, you'll see color flashing as you select windows. Parts of your display may also go black as other color maps are swapped in.

To avoid color flashing, and to conserve what used to be precious color cells, most X applications use the default colormap. Using the default colormap makes your code simpler (you often don't have to do special coding to use this colormap) and helps enforce a unity of colors on your display. You can get the default colormap by calling the `DefaultColormap` macro:

```
Colormap DefaultColormap(
    Display* display,
    int screenNumber)
```

You can create your own colormaps, but here things start to become more complex because you have to deal with visuals.

VISUALS AND COLORMAPS

X abstracts the differences between color hardware through the concept of a visual. X provides six visual abstractions:

- `PseudoColor` visuals provide read-write color arrays. The pixel value provides an index into the array and the contents of the array hold the actual color value (the RGB data). You can change cells in a colormap created under a `PseudoColor` visual.
- `GrayScale` visuals are a lot like `PseudoColor` visuals, except you can only have shades of gray. With such a visual, you must set color cells to hold equal red, green, and blue values.

- `StaticColor` visuals are a lot like `PseudoColor` visuals, but the RGB values in the colormaps are predefined. These colormaps are read-only.
- `StaticGray` visuals are like `StaticColor` visuals, but allow for only predefined shades of gray. A black and white display has a very simple `StaticGray` visual.
- `DirectColor` visuals decompose the pixel values into separate fields for red, green, and blue. Each separate field then provides an index into a separate red, green, or blue colormap.
- `TrueColor` visuals also decompose the pixel values like `DirectColor` visuals. `TrueColor` visuals, though, have read-only RGB values. Typically, these RGB values provide a near linear ramp of colors.

Most of the differences of the visual abstractions relate to the type of colormaps you can create for a given class of visual, such as `PseudoColor` or `DirectColor`.

To create a colormap, call `XCreateColormap`:

```
Colormap XCreateColormap(
    Display* display,
    Window window,
    Visual* visual,
    int allocateFlag)
```

When creating a colormap, you can pre-allocate all the cells in the colormap by passing the flag `AllocAll` as the `allocateFlag` parameter. Pass `AllocNone` to request that no cells be allocated for your application up front. Note that for visual types that only provide read-only colormaps, you must pass `AllocNone`. These visuals are the `StaticColor`, `StaticGray`, and `TrueColor` visual types. It's up to you to remember which visuals lead to which types of colormaps.

WORKING WITH MULTIPLE VISUALS

Like the default colormap, there's also a default visual. Typically, this is a `PseudoColor` visual, since most X applications are coded to assume a `PseudoColor` visual. To get the default visual, call the `DefaultVisual` macro:

```
Visual* DefaultVisual(
    Display* display,
    int screenNumber)
```

To find the visuals available on your display, you can call `XMatchVisualInfo` to find a visual that matches a specific type of visual, such as `DirectColor`, and a preferred depth (number of color planes) that you pass in. If you are not sure which type of visual you want to use or its depth (perhaps you want to find the visual with the greatest depth, for example), then you can call `XGetVisualInfo` to get a list of visuals to choose from. Both of

these functions provide `XVisualInfo` structures, from which you can get the `Visual` pointer you need for the X routines.

The tricky part is that when you create a window, it must be associated with a visual. You can then associate any number of colormaps with a window, but each of these colormaps must have been created for the visual—the same visual the window is associated with. If you don't follow these rules, the X server will generate `BadMatch` errors.

Because of this, it's generally a good idea to use the same visual across all the windows your application creates. Now, this becomes more complicated when you work with X toolkits, such as Motif. Motif tries to hide the creation of actual windows. In most cases, each Motif widget creates a window on the display (windows in X are relatively cheap resources). Motif (actually the underlying X Toolkit Intrinsics library) creates these windows when you call `XtRealizeWidget`, which is often the very last call made before jumping into the event loop.

Furthermore, all Motif widgets have foreground, background, and `borderColor` resources, all of which refer to color cells. These color cells must be valid cells for the colormap associated with the window, which in turn must be created under the visual associated with the window.

In practical terms, you should set the following resources for every top-level shell widget (widgets beneath the top-level shell can inherit the proper values):

- `visual`, which holds the visual you want to use for the window
- `colormap`, which must be a colormap created under the visual
- `depth`, which should match the depth (number of color planes) of the colormap, and be a valid depth supported by the visual
- `background` or `backgroundPixmap`, which must be a valid color cell in the colormap associated with the window or a valid pixmap that has a depth that matches the depth of the colormap
- `borderColor` or `borderPixmap`, which follow the same rules as listed previously
- `foreground`, which must be a valid color cell as described previously

You need to set these resources on each top-level shell widget, before calling `XtRealizeWidget`. Remember that menus and dialogs are also top-level shell widgets. I've found these guidelines helpful:

- Try to use the same visual throughout your application. This eliminates a whole series of `BadMatch` errors where the visuals, colormaps, and color cells don't match up.

- Where possible, use the same colormap throughout the application. High-end imaging applications won't be able to do this, of course. If you cannot use the same colormap, then try to use different colormaps only for special image-processing or other high-color-usage windows. For the rest of the application, use the same colormap. This often works well with a Motif drawing area widget for displaying an image or other color-intensive rendering.
- Always make sure all color cells used in a window come from a colormap associated with that window.

With the rules associated with visuals, colormaps, and color cells, making use of X color capabilities can seem a daunting task, but if you start small and gradually build up your application, it will turn out fine.

REFERENCES

1. From TechWeb's TechEncyclopedia (<http://www.techweb.com/encyclopedia/>), a service of CMPnet and The Computer Language Co. Inc.

THE FREE R.A. BLOCH NATIONAL CANCER HOTLINE **1-800-433-0464**

When cancer strikes you, your family or friends, reliable information is vital. Talking to a cancer survivor will help you or your loved ones successfully cope with this always-treatable, often-curable disease. All at absolutely no charge.

Call **1-800-433-0464**, **24 hours a day** to be matched with a survivor with your type of cancer. Let the R.A. Bloch National Cancer Hotline help you get through the shock and uncertainty of a cancer diagnosis.

All calls are absolutely free, completely sponsored by Richard Bloch and Sprint, and donations are never, ever solicited. Our volunteers don't give medical advice — they talk from personal experience of beating cancer. Their "I've-been-there" experience might save your life.

Call now for help, comfort and hope.

1-800-433-0464

Spotlight

How to get a better high tech job in 3 easy steps:

1. Go to dice.com
2. Go to dice.com
3. Go to dice.com

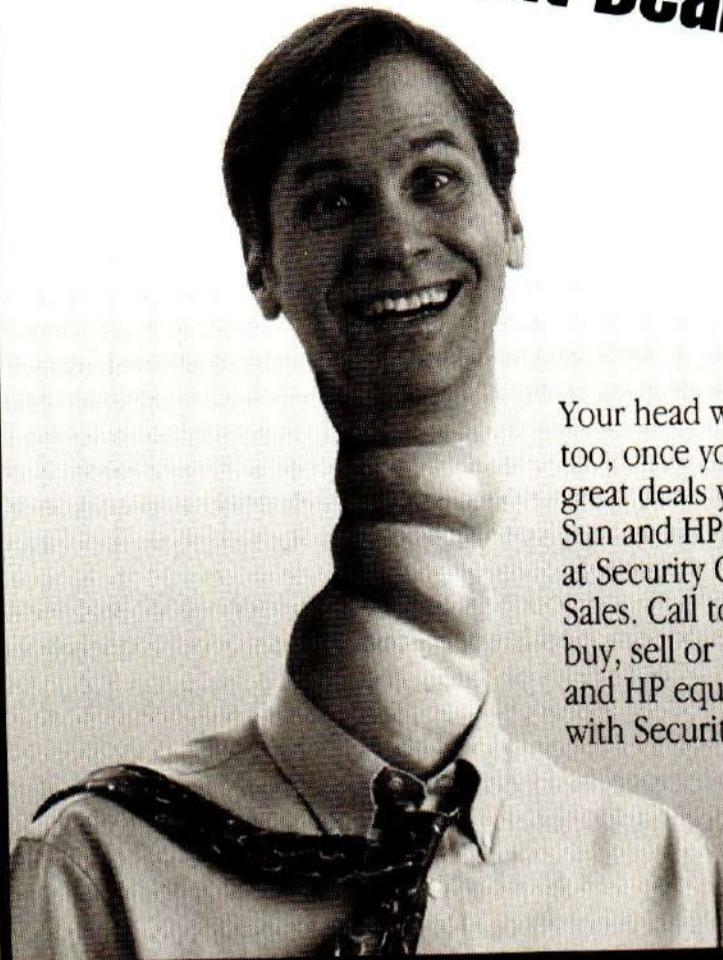


110,000 high tech jobs, including your next one.

AN EARTHWEB SERVICE

Select No. 117 at www.performancecomputing.com/fyi/

Customer's Head Spins Over Great Deals



Your head will spin, too, once you see the great deals you get on Sun and HP equipment at Security Computer Sales. Call today to buy, sell or trade Sun and HP equipment with Security!

SECURITY COMPUTER SALES

Phone: 651.653.5200 FAX: 651.653.4300

E-Mail: Sales@Securitycomputer.com

2340 County Road J • White Bear Lake, MN • 55110

Select No. 174 at www.performancecomputing.com/fyi/



Burning Your Operating Systems at Both Ends?

Let's Face it: integrating UNIX and Windows NT is an immense challenge. You may even be ignoring the migration problem, hoping it will go away. To prevent disaster, you need a resource to guide you through the solution-seeking process.

unixintegration.com is a Web site focused exclusively on helping you integrate and manage your UNIX and Windows NT systems. To help keep you on track as you merge these two complex environments, this site includes a monthly article that focuses on UNIX/NT integration, along with pertinent articles from **Performance Computing** and **Windows NT Systems**. Topics include sharing e-mail resources, scalability, the Internet, and database compatibility.

This site also features an Integration Q&A section that answers pressing technical questions, an Integration Library that guides you to additional resources, and a Product Gallery that gives you an instant connection to pertinent vendor information.

Visit www.unixintegration.com today!

WINDOWS NT
SYSTEMS

unixintegration.com

PERFORMANCE
COMPUTING

Spotlight

EasySpooler

is
ALIVE!

Version five of the most powerful and mature UNIX print management software is now available. If you haven't tried it yet, NOW MAY BE THE TIME! Thousands of EasySpooler customers can't be wrong.

214/522-2324
www.easyspooler.com

Select No. 103 at www.performancecomputing.com/fyi/

I wouldn't store my data on anything else!"

Ride the Wild Surf at
www.partnersdata.com



Partners
Data Systems, Inc.

800.550.3005

No. 141 at www.performancecomputing.com/fyi/

UNIX Pros Shop Here!

SUN, HP, IBM, SCO

	LIST	OURS
Cheyenne ArcServe/Open Backup(SCO=\$550)	1995	1240
VSI FAX Gold Fax Server, 5 Win Client (SCO=\$995)	2695	2225
Term Comm./File Transfer(SCO=\$390)	695	595
Legato Networker Backup S/W(SCO=\$1275)	2500	2050
DoubleVision Remote Supp. S/W(SCO=\$330)	895	730
WordPerfect 8.0 1st user CD	495	395
FineLine Message Manager E-Mail/Apps 25 mailbox	800	710
FrameMaker 1-user Floating License (no SCO)	2595	2285
Uniplex II Plus 20-user	5500	4895
Informix On-Line Dynamic Server 1-user (min. 10)	1500	Call
Informix On-Line Workgroup Server 1-user (min 5)	295	Call
Seagate Cheetah 18.2GB U2-SCSI External	1495	1260
Yamaha CD-RW recorder Ext. 2x/4x/6x	550	415
HP Surestore 24GB 4mm DAT Tape Drive External	1499	1185
HP SureStore 70GB DLT Tape Drive External	7188	6150
HP ProCurve Hub 10/100BT 24 ports w/SNMP	1199	1020

PC to UNIX Connectivity

Hummingbird eXceed PC X-Server NT/95	545	399
Hummingbird Maestro NFS Server WIN NT	599	420
FacetWIN emul, file/print/modem share 10-user	1600	1350
ICE.TCP.Pro (DOS/Win -> UNIX, 10 PC's)	895	695
Tiny Term Plus w/ TCP 5-user pack	495	375

**Complete x86 Systems (HP NetServers • UNIX, LINUX, NT)
Plus 1000's of other products not listed here!**

Other Platforms: Solaris x86, UnixWare, WinNT, Linux, DEC, DG, SGI

Open Systems

EXPRESS!

Your direct source for UNIX,
Win NT & networking products

46 Whitney Road, Suite B8, Mahwah, NJ 07430

1-800-445-9292

Search and order online at
www.osexpress.com

E-Mail us at: sales@osexpress.com MC/VISA/AMEX/DISCV, Tel: 201-847-8880
C.O.D., & Net Terms Fax: 201-847-8899

Select No. 111 at www.performancecomputing.com/fyi/



Spatch

wireless text-messaging software
for Unix & NT

Spatch sends full text-messages to
pagers and phones

Send messages from:

- the command line
- the user interface
- your intranet
- e-mail
- other applications



**Starting at
only \$199**

The Hyde Company
Duluth, Georgia
spatch@hydeco.com

770-495-0718

www.spatch.com

Select No. 177 at www.performancecomputing.com/fyi/



Get in Touch with Trident.

TOUCH SCREEN TECHNOLOGY EXCELS AT

- Simplifying computer input
- Saving space
- Supporting fast computer operations
- Reducing maintenance costs and down-time

AND NOW, YOU CAN HAVE TOUCH WITHOUT MODIFYING YOUR EXISTING APPLICATIONS. TRIDENT'S TOUCH SOLUTIONS INCLUDE

- Touch-configured monitors, X terminals, and Netstations
- Touch screen integration with your monitors
- Touch screen hardware and driver software
- All major touch technologies
- Rack mounting or ruggedized NEMA enclosures and kiosks with touch
- Touch application software development services

TRIDENT'S TOUCH SOLUTIONS WORK WITH

UNIX Workstations Netstations

DOS & Windows UNIX PC's

X Terminals Macintosh

Trident puts you in touch.™

TRIDENT SYSTEMS INC.
TEL 703.273.1012
FAX 703.273.3763

email: touch@tridsys.com
url: <http://www.tridsys.com/cpg>

Select No. 114 at www.performancecomputing.com/fyi/

**Get Your
Reprints
Here!**

Copies of product reviews, articles, or columns are available through our reprints department. If your company was mentioned in *Performance Computing*, you need to make sure that the right people see it!

**Order Your Reprints Through
Andrea Ramiza at
(415) 905-2591**

**PERFORMANCE
COMPUTING**

Spotlight

Do You Have Problems With UNIX Or NT Printing?

Save time and money by controlling print management operations!



Lets You:

- Intuitively control, manage & view printers & jobs (with browser, X-GUI, or command line interface)
- Restart print jobs at any page
- Re-direct jobs to other printers
- View jobs before or without printing
- Wait to purge files from queue after printing
- Print page ranges or even a single page
- Integrate with ERP applications (SAP, Oracle, etc.)
- Intelligent Report Distribution
- Tracks printer usage for charge backs
- Share Windows and UNIX printers using WinPrint
- Associate jobs with specific forms

Free trial version on the web!
www.plustechologies.com



Sales/Support: **TOLL FREE** (877) 899-PLUS
International: (513) 746-8118

Select No. 121 at www.performancecomputing.com/fyi/



Customer Goes Ape Over Great Deals

You'll go ape, too, once you see the great deals you get on equipment at Security Computer Sales. Whether you're buying, selling or trading Sun and HP equipment, go with Security! We don't monkey around. Call us today and see for yourself.

SECURITY COMPUTER SALES

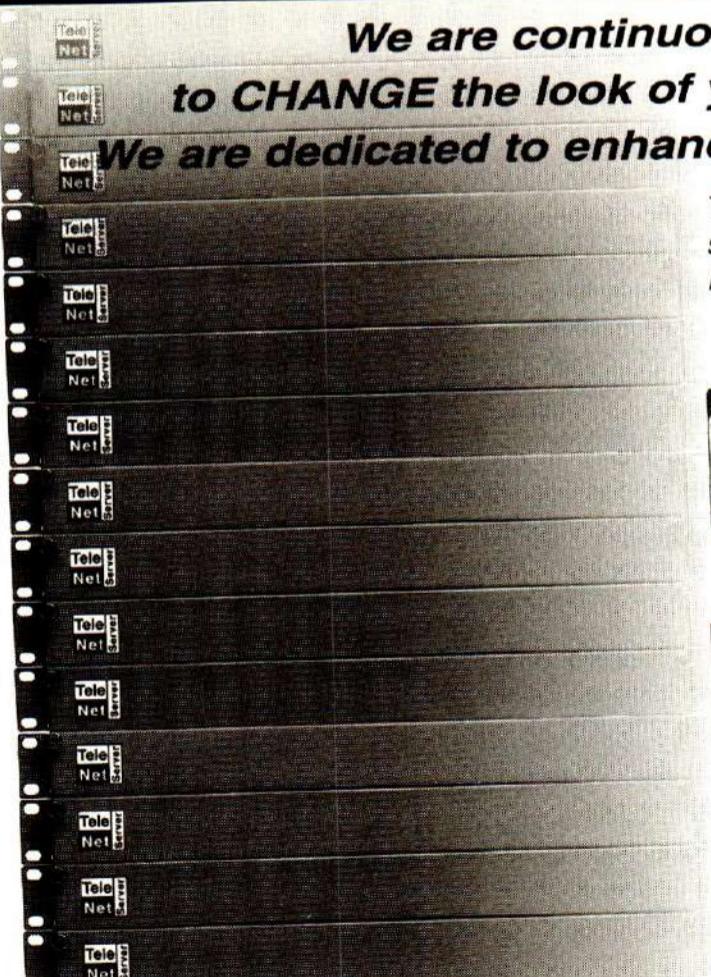
Phone: 651.653.5200 FAX: 651.653.4300

E-Mail: Sales@Securitycomputer.com

2340 County Road J • White Bear Lake, MN • 55110

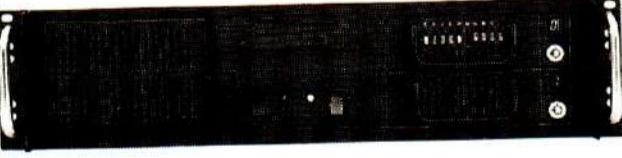
Select No. 174 at www.performancecomputing.com/fyi/

We are continuously working to CHANGE the look of your DATACENTRE. We are dedicated to enhance the NET ECONOMY.



Telenet offers state-of-the-art, cutting edge; internet servers (Intel x86 and Sun Sparc), RAID solutions and Multiboot-Workstation products at very competitive prices.

CompactRack Server
2.5U based configurations



SlimRack Server
2U based configurations



MicroRack Server
1U based configurations



Many other configurations available. Please visit our website for more information.
Custom configurations • RAID Servers • Internet Servers • OS of your Choice

Telenet
System Solutions, Inc.

1-877-7Telenet www.tesys.com
2480 Kruse Drive, San Jose, CA 95131 • TEL (408) 383-0334 • FAX (408) 383-0335

Select No. 176 at www.performancecomputing.com/fyi/

PERFORMANCE COMPUTING

SUBSCRIBE ONLINE @
<http://www.performance-computing.com>

FREE SUBSCRIPTION ORDER CARD

Note: This card may also be used to: 1. Renew your subscription
 2. Change your address. For change of address, OLD ADDRESS
 MUST BE SHOWN. Attach label from recent issue or write old
 address, and provide all information requested on this card.

YES! I want to receive/continue to receive *Performance Computing* magazine **FREE!**
 No

Signature **X**

(required)

Name _____

Date _____

Company _____

Title _____

Mailing Address _____

City _____

State _____

Zip _____

Phone () _____

Fax () _____

Email _____

We respect the privacy of your email address. Please indicate those parties you **DO NOT** wish to contact you via email:
 Performance Computing Magazine Other Miller Freeman, Inc. products and services Relevant outside business organizations and vendors that may be of interest to you.

IMPORTANT: ALL questions must be answered to start subscription.

1 Are you involved in specifying, recommending, purchasing and/or approving computer systems, software, and/or services for use within a data center environment? (Data center is defined as any location where high performance computer systems are managed for access by corporate and/or Internet users)

01 YES 02 No

2 Which of the following BEST describes your job title? (Check ONE only)

IS/Computer Systems Operation

- 01 Management: VP IS, Director/Manager IS, Director/Manager or Systems and/or Network Administration, Manager/Director of Operations, Data Center Manager, Director/Manager of Applications Development/Programming
- 02 Supervisory: Supervisor Equipment/Analysis, Chief/Lead/Senior Project Leader, Senior Systems Analyst, Senior Systems/Network Administrator
- 03 Staff: Systems Analyst, Software Engineer, System/Network Administrator, Specialist, Software Developer, Technical Staff

Corporate Management

- 04 General/Corporate: President/Owner/Director/Chairman/Partner, Vice President/General Manager
- 05 Financial/Administrative: Comptroller/Treasurer, VP Finance, Chief Accountant/CFO, VP Operations, Director or Manager of Purchasing

Consultant

- 06 Consultant to IT Organization
- 10 Other (please specify) _____

3 Which of the following job functions are you involved with? (Check ALL that apply)

- 01 Systems Management/Administration
- 02 Network Management/Administration
- 03 Internet/Intranet Management/Administration
- 04 Telecommunications Management
- 05 Database Management/Administration
- 06 Applications Development/Integration
- 19 Other (please specify) _____
- 20 None of the above

4 Which of the following products do you specify, select, or approve? (Check ALL that apply)

- | | |
|---|---|
| 01 <input type="checkbox"/> Operating Systems | 14 <input type="checkbox"/> Hard Disks/Solid State Disks |
| 02 <input type="checkbox"/> Midrange Computers | 15 <input type="checkbox"/> Bridges/Routers/Switches/Hubs |
| 03 <input type="checkbox"/> Workstations | 16 <input type="checkbox"/> Printers |
| 04 <input type="checkbox"/> Personal Computers | 17 <input type="checkbox"/> UPS |
| 05 <input type="checkbox"/> Network Computers | 18 <input type="checkbox"/> Software Development Tools |
| 06 <input type="checkbox"/> Database/File Servers | 19 <input type="checkbox"/> Web Development Tools |
| 07 <input type="checkbox"/> Applications Servers | 20 <input type="checkbox"/> Java Tools |
| 08 <input type="checkbox"/> Internet/Intranet Servers | 21 <input type="checkbox"/> Systems Management Software |
| 09 <input type="checkbox"/> Terminal Servers | 22 <input type="checkbox"/> Network Management Software |
| 10 <input type="checkbox"/> Mass Storage Systems | 23 <input type="checkbox"/> Security/Firewalls |
| 11 <input type="checkbox"/> RAID Storage | 24 <input type="checkbox"/> Connectivity Software |
| 12 <input type="checkbox"/> Tape Libraries/Drives | 25 <input type="checkbox"/> Consulting Services |
| 13 <input type="checkbox"/> Backup Software | 26 <input type="checkbox"/> Training Services |
| 40 <input type="checkbox"/> None of the above | |

5 Which of the following BEST describes your organization's primary business activity? (Check ONE only)

- 01 Manufacturing/Production (non-computer)
- 02 Government
- 03 Educational
- 04 Finance
- 05 Trade/Wholesale/Retail
- 06 ISP (Internet Service Provider)
- 07 Communication Service Provider/Telecom
- 08 Medical and legal services
- 09 Transportation services
- 10 Utilities
- 11 Agriculture/Construction/Mining
- 12 Computer Data Processing Services
- 13 Media/Advertising/Web
- 14 VAR (value added reseller of hardware and/or software)
- 15 ISV (independent software vendors)
- 16 Manufacturer of computer systems or related hardware
- 17 Systems/Network Integration
- 18 Independent Consulting Firm
- 30 Other (please specify)

6 What is the annual expenditure (or sales if you are a reseller) for computer products and services within your organization?

- 01 Over \$10 Million
- 02 \$5 Million - \$9.9 Million
- 03 \$1 Million - \$4.9 Million
- 04 \$500,001 - \$999,999
- 05 \$100,001 - \$500,000
- 06 \$50,001 - \$100,000
- 07 Less than \$50,000

7 What is the annual expenditure (or sales) for computer products and services that you personally select, approve, or influence? (Check ONE only)

- 01 Over \$10 Million
- 02 \$5 Million - \$9.9 Million
- 03 \$1 Million - \$4.9 Million
- 04 \$500,001 - \$999,999
- 05 \$100,001 - \$500,000
- 06 \$50,001 - \$100,000
- 07 Less than \$50,000

8 What is the total number of servers and clients installed in the entire enterprise within which you work and/or consult for? (Check ONE only)

- | | |
|--|--|
| 01 <input type="checkbox"/> 10,000 or more | 05 <input type="checkbox"/> 100 - 499 |
| 02 <input type="checkbox"/> 5,000 - 9,999 | 06 <input type="checkbox"/> 10 - 99 |
| 03 <input type="checkbox"/> 1,000 - 4,999 | 07 <input type="checkbox"/> Less than 10 |
| 04 <input type="checkbox"/> 500 - 999 | |

9 Your organization uses/resells/develops products running on the following operating systems: (Check ALL that apply)

- | | | |
|--|---|---|
| 01 <input type="checkbox"/> AIX | 07 <input type="checkbox"/> Linux | 13 <input type="checkbox"/> Windows 2000/Datacenter |
| 02 <input type="checkbox"/> Digital Unix | 08 <input type="checkbox"/> OS/400 | 14 <input type="checkbox"/> OS/2 |
| 03 <input type="checkbox"/> HP/UX | 09 <input type="checkbox"/> MVS | 15 <input type="checkbox"/> NetWare |
| 04 <input type="checkbox"/> Solaris | 10 <input type="checkbox"/> Windows 3.x/95 | 16 <input type="checkbox"/> Mac OS |
| 05 <input type="checkbox"/> SGI/IRIX | 11 <input type="checkbox"/> Windows NT | 17 <input type="checkbox"/> DOS |
| 06 <input type="checkbox"/> SCO UNIX | 12 <input type="checkbox"/> Windows 2000/Server | 18 <input type="checkbox"/> VMS |
| 30 <input type="checkbox"/> Other (please specify) | | |

Fax both sides to 847-647-5972

continued on back →

49LQC

10 What operating system platforms do you personally use/manage/or resell now or plan to within the next 12 months? (Check ALL that apply)

- 01 UNIX
- 02 Windows NT/2000
- 03 MVS
- 04 VMS
- 05 OS/400
- 06 Linux
- 30 Other (please specify) _____

11 What is the total storage capacity of the data center(s), in gigabytes, you work within and/or manage? (Check ONE only)

- 01 50,000 or more
- 05 500 - 999
- 02 10,000 - 49,999
- 06 100 - 499
- 03 5,000 - 9,999
- 07 10 - 99
- 04 1,000 - 4,999
- 08 Less than 10

12 What is the average total storage capacity per server, in gigabytes, of the data center(s) you work within and/or manage?

- 01 10,000 or more
- 05 100 - 499
- 02 5,000 - 9,999
- 06 10 - 99
- 03 1,000 - 4,999
- 07 Less than 10
- 04 500 - 999

13 What is the average total universe of users accessing your data center(s) daily?

- 01 Over 1 Million
- 05 10,000 - 49,999
- 02 500,000 - 1 Million
- 06 1,000 - 9,999
- 03 100,000 - 499,999
- 07 Less than 1,000
- 04 50,000 - 99,999

14 What is the highest level of the enterprise for which you are involved in recommending, specifying, or approving purchases of hardware, software, and/or services? (Check ONE only)

- 01 More than one company
- 05 Single department
- 02 Entire company
- 06 Yourself only
- 03 Multiple divisions or sites
- 20 Other (please specify) _____
- 04 Multiple departments

15 Which of the following enterprise software products are installed within your organization or do you resell? (Check ALL that apply)

- 01 Oracle
- 07 CA Unicenter
- 02 Sybase
- 08 HP Openview
- 03 Informix
- 09 Tivoli
- 04 IBM DB2
- 10 SMS Server
- 05 MS SQL Server
- 11 Internet Information Server
- 06 SNA Server
- 30 Other (please specify) _____

Publisher reserves the right to determine qualification for free subscriptions.

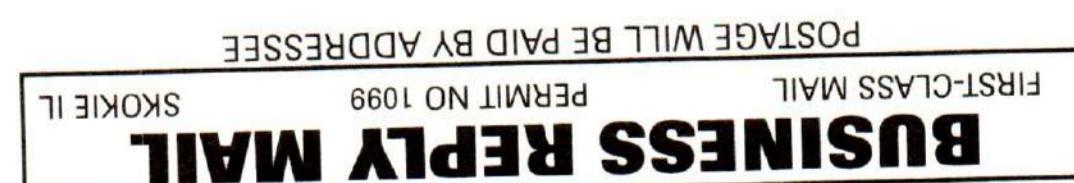
Please send *Performance Computing* magazine to the following employees at this location:

1. Name	Title
2. Name	Title
3. Name	Title
4. Name	Title
5. Name	Title

For fastest service
fax both sides to
847-647-5972

SKOKIE IL 60076-9430
PO BOX 1280

COMPUTING
PERFORMANCE



UNITED STATES
IN THE
IF MAILED
NECESSARY
NO POSTAGE



Please seal with tape along this edge only. Thank you.

INDEX OF ADVERTISERS

Request product information online with the fastest, most convenient way to contact our advertisers.



How to use FYI:

1. Go to www.performance-computing.com
2. Click on the FYI link
3. Search by product category, product name, company name, or by the FYI number listed below
4. Choose how you wish to receive information (mail, e-mail, or phone) or link directly to vendor Web site
5. Fill out contact information

Advertisers In This Issue

	PAGE	FYI		PAGE	FYI
ADIC/EMASS	22	137	KL Group	21	7
Century Software	57	21	Open Systems Express	63	111
Computer Associates	11,C4	3	Partners Data System Inc.	63	141
Data Focus	2	4	Radiant Resources	25	171
Dice /D&L Online	62	117	Rave Computers	14	39
Digital Controls Corporation	64	121	Resillience Corp.	19	178
Easy spooler	63	103	Security Computer Sales	62,64	174
Graphon	16	138	SGI	12	10
Hyde Co.	63	177	Software Engineering of America	C2	190
IBM	4	52	Telenet	64	176
Integrix	15	140	Trident Systems Inc.	63	114
Intel	6		VA Linux	31	183
J. River	36	2	Veritas	2	93
JSB Corporation	26	99	Winchester Systems Inc.	C4	85
Kingston	8	18			

The index on this page is provided as a service to readers. The publisher assumes no liability for errors or omissions.

PERFORMANCE COMPUTING

411 Borel Ave., Suite 100
San Mateo, CA 94402
Tel: 650-358-9500
Fax: 650-358-9739
<http://www.performance-computing.com>

PUBLISHER
John R. Keough
jkeough@mfi.com

VICE PRESIDENT/CIRCULATION
Jerry M. Okabe

MARKETING MANAGER
Katherine Hartlove

CIRCULATION MANAGER
Lisa Eversole
leversole@mfi.com

NEWSSTAND MANAGER
Joyce Gorsuch

ADVERTISING PRODUCTION COORDINATOR
Martin Ruiz 650-655-4383

REPRINTS
Andrea Ramiza 415-905-2591

SALES OFFICE-EAST

One Penn Plaza
New York, NY 10119

EASTERN SALES MANAGER
Anthony Spizzirro
212 615-2743/aspizzirro@mfi.com

SALES OFFICE-WEST

411 Borel Ave., Suite 100
San Mateo, CA 94402

NATIONAL SALES MANAGER
Ken McGlothlin
650-655-4298/kmcglothlin@mfi.com

SALES OFFICE-U.K./EUROPE/MIDDLE EAST

Brian Osborn
Breakout Marketing
+49-431-801740
breakout_mktg@compuserve.com

ONLINE ADVERTISING SALES MANAGER
Roddy DeLaGarza
415-905-2351/rdelagarza@mfi.com

SUBSCRIPTION INFORMATION
888-847-6188
847-647-8603
performancecomputing@halldata.com

Copyright 1999 by Miller Freeman Inc. Rights reserved. Contents may not be reproduced without permission. Performance Computing is a trademark of Miller Freeman Inc.

MILLER FREEMAN
A United News & Media publication
V BPA

Bottled at the Source



Code Solutions so Pure They Must be from Microsoft®!

SAVE OVER 40% NOW!



Subscribe Today

*M*icrosoft Systems Journal is the leading source of in-depth technical information for developers taking advantage of the Windows family of operating systems.

*W*ritten by industry authorities and edited for technical accuracy by programmers at Microsoft, *MSJ* is the only publication on Windows-based development with this unique access to the designers of today's hottest technologies.

A must-read for any developer targeting Windows!

**YES! Start my subscription to
MICROSOFT SYSTEMS JOURNAL for only
\$34.95! I get a full year's worth—12 issues
in all—over 40% off the cover price!
That's a savings of almost \$25!**

NAME _____

COMPANY _____

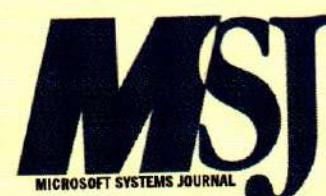
STREET ADDRESS _____

CITY/STATE/ZIP _____

EMAIL ADDRESS (REQUIRED) _____

5HP8

Call Today: **Or FAX:**
800-666-1084 **303-661-1885**
Or Mail to:
Microsoft Systems Journal
P.O. Box 56621
Boulder, CO 80322-6621



www.microsoft.com/msj

STORAGE ATL Products

ATL Products introduced three Fibre Channel solutions as part of the company's Network Storage Systems strategy for Storage Area Network (SAN) connectivity. The first of ATL's Fibre Channel solutions is an embedded Fibre Channel interface option that is factory-installed for new purchases of the P1000 Series and P3000 Series libraries, which are based on the company's Prism Library Architecture. For systems already deployed in

the field, ATL's second Prism-based solution consists of a Fibre Channel bridge field-upgrade kit that converts the SCSI-based library into a SAN-enabled library. For those who currently use or are planning to use non-Prism-based ATL autoloaders and libraries, ATL features a Crossroads 4200 external Fibre Channel bridge. This solution lets all ATL SCSI-based libraries be connected to Fibre Channel SANs.

The P1000 and P3000 Series embedded solutions are designed to connect to Fibre Channel loop or switched-fabric topologies. They support standard multimode (short wave) optical-fiber cabling, and provide cable lengths up to 500m between the library and host computer (or a hub or switch). When de-



ATL's P3000

ployed with Fibre Channel, the P1000 and P3000 will connect directly to a SAN to back up multiple servers simultaneously.

The kit is designed to be field-installed within the P1000 or P3000 library and can connect these libraries to a Fibre Channel network via multimode optical cabling.

For ATL legacy libraries, the Crossroads 4200 Fibre Channel bridge connects two differential SCSI buses to one Fibre Channel connection using multimode optical cabling.

The P1000 Series is designed for departmental and midrange storage requirements, featuring up to four DLTtape drives and 30 cartridges for 72GB/hour performance and 1.05TB of storage capacity. The P1000 also

features hot-swappable drives and user-replaceable components.

The P3000 Series is a DLTtape library with a high-availability (HA) design similar to HA servers and RAID systems. Targeted for enterprise-class applications, the P3000 features up to 16 DLTtape drives and 326 cartridges for 288GB/hour performance and 11.4TB of storage capacity. Its HA design features single-connector hot-swappable DLT-tape drives, power supplies, and fans, as well as redundant power supplies, fans, and AC sources. The hot-swappable components allow for non-disruptive maintenance, while the redundant AC components allow online replacement and continual uptime if a single component fails.

Contact ATL Products, 2801 Kelvin Ave., Irvine, CA 92614-5872. Tel.: 949-856-7800, fax: 949-774-6900, e-mail: atlinfo@atlp.com, Web: <http://www.atlp.com>.

FYI No. 152

MTI Technology Corp.

MTI Technology Corp. released the NorthStar series, a line of products that combine Fibre Channel switched-fabric technology, high-availability (HA) information servers, and integrated management services.

The NorthStar system is a fault-resilient, scalable storage system incorporating MTI's HA Gladiator 6700 RAID technology that lets users expand storage to 3TB of data. The system offers multiple concurrent host support for applications that require high bandwidth and transfer rates, such as large enterprise data warehousing, data mining, Internet/intranet/extranet applications, and online transaction processing (OLTP) environments.

NorthStar includes MTI's DataAssure technology, which protects host/array data integrity, and Write-

Guard, an MTI process that ensures writes in progress are completed in the event of a power interruption. NorthStar also features a built-in management processor that provides a GUI-based storage-management facility. This feature enables centralized administration of all aspects of storage configuration, fault notification, and performance monitoring of one or more MTI storage servers from anywhere on the network, anywhere in the world.

The NorthStar series is supported by AIX, HP-UX, IBM, IRIX, NetWare, NT, and Solaris. The system allows capacity scalability with disk expansion to 3TB in a single cabinet.

Entry-level system pricing starts at \$129,000 and ranges to \$400,000 for a fully configured system.

Contact MTI Technology Corp., 4905 E. La Palma Ave., Anaheim, CA 92807. Tel.: 714-970-0300 or 800-999-9MTI, fax: 714-693-2256, e-mail: info@mti.com, Web: <http://www.mti.com>.

FYI No. 155

Logic Storage Systems Inc.

The MetaStor Storage Solutions Division of LSI Logic Storage Systems Inc. announced support of 50GB 7,200-rpm disk drives for its MetaStor line of scalable storage systems. The 50GB drives can hold 5TB in a 72-inch cabinet, and are SCSI and Fibre Channel compatible. The drives extend system flexibility by enabling a capacity of up to 5TB in a single 72-inch cabinet.

The 50GB drive costs \$3,995. The 9GB 10,000-rpm and 18GB 10,000-rpm drives list for \$1,395 and \$2,750, respectively.

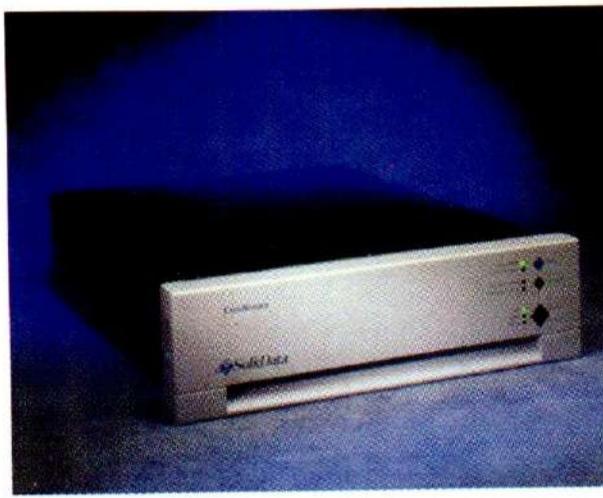
Contact MetaStor, 1551 McCarthy Blvd., Milpitas, CA 95035. Tel.: 888-METASTOR, e-mail: info@metastor.com, Web: <http://www.metastor.com>.

FYI No. 192

Solid Data Systems Inc.

Solid Data Systems Inc. introduced Excellerator 800 FC, a solid state storage system utilizing a Fibre

Channel interface for host connection. The system supports UNIX and Windows NT environments, and lets storage area networks reach their maximum performance by providing high-speed storage that can be allocated to the applications that have the most pressing performance requirements.



SDS's Excellerator 800 FC

Excellerator 800 FC lets customers improve I/O-intensive applications such as e-commerce, Internet, e-mail and news, trading, and customer care. The system incorporates a 100MB/sec Fibre Channel host connection that supports the connectivity, bandwidth, and cable-length distances required for server clusters and Fibre Channel Arbitrated Loop (FC-AL) configurations. The system also supports SAN hub and switched fabric configurations.

Solid Data's Excellerator 800 FC is a self-contained storage system utilizing DRAM architecture to increase the performance of I/O-intensive applications running in online, real-time environments.

Excellerator 800 FC supports capacities up to 5.4GB. In addition, Excellerator 800 FC offers a dual-port option to enable high-availability failover in mission-critical environments.

The Excellerator 800 FC is priced at \$27,225 for a 536MB system and \$80,550 for a 5.4GB configuration.

Contact Solid Data Systems Inc., 2945 Oakmead Village Ct., Santa Clara, CA 95051. Tel.: 408-727-5497, fax: 408-727-5496, Web: <http://www.soliddata.com>.

FYI No. 191

NETWORKING

Allied Telesyn Intl.

Allied Telesyn Intl. released a line of 100Mbps long-haul media converters, the rackmountable AT-MC103LH and AT-MC104LH, that make copper to fiber or multimode fiber to single-mode fiber connections across distances of up to 40km.

The AT-MC103LH media converter converts 100Base-TX to 100Base-FX for single-mode fiber with an SC connector; the AT-MC104LH converts 100Base-FX for multimode fiber with an SC connector to 100Base-FX single-mode with an SC connector; in each case extending distance limits to 40km.

Both media converters have Allied Telesyn's MissingLink feature, and a fiber test switch for link checks. The converters also have a user-selectable switch for half- or full-duplex that lets them integrate Fast Ethernet switches and hubs into the same network. Each media converter unit is equipped with diagnostic LEDs, and both the AT-MC103LH and AT-MC104LH can be used in a standalone or rackmount environment.

The AT-MC103LH and AT-MC104LH are priced at \$1,010 and \$1,106, respectively, carry a lifetime warranty, and come with free, unlimited technical support.

Contact Allied Telesyn Intl., 19015 N. Creek Pkwy., Suite 200, Bothell, WA 98011. Tel.: 425-487-8880 or 800-424-4284, fax: 425-489-9191, Web: <http://www.alliedtelesyn.com>.

FYI No. 144

SOFTWARE MANAGEMENT

NasTel Technologies

NasTel Technologies began shipping the latest release of AutoPilot/ORB, a middleware management solution designed to provide a unified application view of multiple underlying middleware systems from a centralized point.

SD BOOTCAMP SERIES

Add new weapons to your software development arsenal and learn the skills and discipline necessary to optimize your COM or Java programming.

ENLIST IN 5 DAYS
OF INTENSIVE
HANDS-ON TRAINING

COM BOOTCAMP

- ★ Boston, Aug. 2-6
- ★ Chicago, Aug. 16-20
- ★ Dallas, Aug. 23-27
- ★ Bellevue, WA, Aug. 30 - Sept. 3
- ★ Santa Clara, CA, Sept. 27 - Oct. 1
- ★ New York City, Oct. 25-29

WHAT YOU'LL EXPERIENCE:

- ★ Master core techniques for building COM servers, objects and components through hands-on experience
- ★ Learn to test from clients written in MFC and ATL
- ★ Keep pace with continually evolving COM extensions such as COM+ & DCOM

EXPERT INSTRUCTOR

Richard Hale Shaw

As an authority on COM & Windows programming with C++, Shaw has written for *PC Magazine*, *Microsoft Systems Journal*, & *VC++ Developer's Journal*. He's also the COM Track Chair for SD'99. www.richardhaleshaw.com

JAVA BOOTCAMP

- ★ New York City, Aug. 9-13
- ★ Dallas, Aug. 23-27
- ★ Bellevue, WA, Aug. 30 - Sept. 3
- ★ Santa Clara, CA, Sept. 27 - Oct. 1
- ★ Chicago, Oct. 4-8
- ★ Research Triangle Park, NC, Oct. 18-22

WHAT YOU'LL EXPERIENCE:

- ★ Dive into Java programming, from basic constructs and syntax to design & application
- ★ Explore advanced topics such as multi-threading, networking, applets, inner-classes, and GUI-building
- ★ Reduce your development time and boost your productivity with Java

EXPERT INSTRUCTOR

Paul Tyma

Tyma is president of preEmptive Solutions, Inc. a Java training, consulting, and products company, and is the lead author of the Waite Group's Java Primer Plus.

Contact us today!

w www.sdexpo.com/bootcamp
f 415-905-2222

Miller Freeman

The Bootcamp Series is brought to you by Miller Freeman Inc., producer of SD'99, the largest independent conference and expo dedicated to developers.

t 800-441-8826
e bootcamp@mfi.com



With AutoPilot/ORB, users collect information about Iona's Object Request Broker (ORB) and publish it back to managers so appropriate actions can be taken. AutoPilot lets managers automate procedures and monitoring conditions based on their specific business requirements. It eliminates the need to write automation procedures and scripts. AutoPilot/ORB lets users monitor, manage, and administer Iona's ORB in conjunction with other related components within, and provides a single point of control for CORBA. As a result, there is a coherent and integrated approach to the management of multiple middleware technologies within an enterprise.

AutoPilot/ORB is platform independent. Pricing for AutoPilot/ORB starts at \$20,000.

Contact NasTel Technologies Inc., 200 Old Country Rd., Mineola, NY 11501. Tel.: 516-294-6033, Web: <http://www.nastel.com>.

FYI No. 197

SERVERS

Compaq Computer Corp.

Compaq Computer Corp. is shipping bundles of Compaq ProLiant Cluster servers, coupled with StorageWorks RAID Array 4000 and 8000 storage subsystems, and cluster software kits.

The bundles are sized for workgroups, departments, and business-critical enterprise environments. Factory-configured and tested by Compaq Custom Systems, the bundles are designed to work with many applications including Microsoft Exchange, Lotus Notes, and file and print services for customers.

The pre-packaged bundles include two dual-processor ProLiant 1850R servers, each with a 450MHz CPU and 128MB RAM, and the StorageWorks RAID Array 4000 with 50GB of hot-plug storage. These support Windows NT or Novell NetWare and are priced at \$49,995 to \$219,995.



Compaq Computer Corp.'s ProLiant Cluster Server

Contact Compaq Computer Corp., P.O. Box 692000, Houston, TX 77269-2000. Tel.: 281-370-0670, fax: 281-514-1740, Web: <http://www.compaq.com>.

FYI No. 196

SOFTWARE CONNECTIVITY

Wall Data Inc.

Wall Data Inc. announced the availability of RUMBA 2000, Web-to-Host version 3.0. RUMBA 2000, Web-to-Host can be installed on most Web servers and offers enterprise organizations the capability of hands-free deployment to clients requiring immediate host access.

Business users are able to work with information residing on IBM Mainframe, IBM AS/400, Hewlett-Packard, UNIX, and DEC VAX systems. For system managers, the product provides streamlined administration and deploys within minutes in any Web-enabled environment.

The latest version of RUMBA 2000, Web-to-Host includes a drag-and-drop browser-based FTP client, as well as support for nearly all in-

dustry-standard application program interfaces (APIs). Organizations can give users browser-based access to existing 16-bit or 32-bit HLLAPI-based applications without re-coding.

RUMBA 2000, Web-to-Host offers scalable deployment options to match an organization's network requirements. The options include two sizes of ActiveX controls (3MB or 300KB) as well as one size of JavaBean (350KB). For custom application development, the product provides ActiveX and JavaBeans development architectures, and centralized software deployment. RUMBA 2000, Web-to-Host also provides a common user interface for all host types, reducing user training requirements and simplifying IT administration.

English versions of RUMBA 2000, Web-to-Host, v. 3.0 cost \$150 per seat. Site licensing and volume discounts are also available.

Contact Wall Data Inc., 11332 N.E. 122nd Wy., Kirkland, WA 98034-6931. Tel.: 425-814-9255 or 800-755-9255, Web: <http://www.walldata.com>.

FYI No. 198

SECURITY

Freshwater Software

Freshwater Software released SiteScope 4.5 performance-monitoring software that can monitor and verify encrypted financial transactions.

SiteScope 4.5 features include a URL Transaction Wizard that helps system administrators create e-commerce transaction monitors. Version 4.5 also supports remote monitoring of additional UNIX platforms, including HP-UX and Linux, and multiple login accounts that let system administrators define different access permissions for different groups (both internal and external).

The URL transaction monitor verifies Web application performance, ensuring that back-end databases are executing properly. Administrators can access SiteScope over the

Internet and correct errors remotely via its HTML interface. SiteScope can manage heterogeneous NT/UNIX environments from one central console.

SiteScope runs on IRIX, Solaris, and NT.

Licenses begin at \$495 for NT and \$1,295 for UNIX. Volume discounts, data-center licenses, and educational pricing are available.

Contact Freshwater Software Inc., 1965 North 57th Ct. Suite 204, Boulder, CO 80301. Tel.: 303-443-2266, fax: 303-545-9533, e-mail: info@freshtech.com, Web: <http://www.freshtech.com>.

FYI No. 194

VeriSign Inc.

VeriSign Inc. released Go Secure!SM for Web Applications, a fast-track approach for enterprises deploying secure business-to-business and business-to-consumer e-commerce applications. The product lets you incorporate digital certificates without requiring proprietary client software or hardware.

Several components, including automation of certificate acquisition, user authentication, and password replacement for access control, reduce the time necessary to set up and deploy these applications as well as simplify end users' acquisition and use of certificates within the application.

VeriSign's Go Secure! for Web Applications service includes a Personal Trust Agent which lets users acquire certificates and transparently use certificates within Web applications; a Password Migration service which allows for automatic replacement of passwords for digital certificates in a highly scalable service; and Administrative Access Control components which help create and set up secure Web applications through canned, preconfigured Web front-end pages.

Go Secure! for Web Applications is included at no charge with bundled OnSite configurations.

Contact VeriSign Inc. 1350 Charleston Rd., Mountain View, CA

94043. Tel.: 650-961-7500, fax: 650-961-7300, Web: <http://www.verisign.com>.

FYI No. 193

SYSTEMS MANAGEMENT

Datametrics Systems Corp.

Datametrics Systems Corp. announced the release of ViewPoint performance data collector for HP-UX 11 systems. This release adds support for both 32-bit and 64-bit computing environments.

The ViewPoint product suite lets system administrators, DBAs, and IT managers take more control of their system's performance across a distributed heterogeneous environment.

ViewPoint can analyze thousands of data points in seconds and make instant correlations from a centralized Windows-based workstation.

Pricing for the ViewPoint data collector for HP-UX starts at \$4,500, depending on the system configuration. It also runs on several other platforms.

Contact Datametrics Systems Corp., 12150 E. Monument Dr. #300, Fairfax, VA 22033. Tel.: 888-436-6300, Web: <http://www.datametrics.com>.

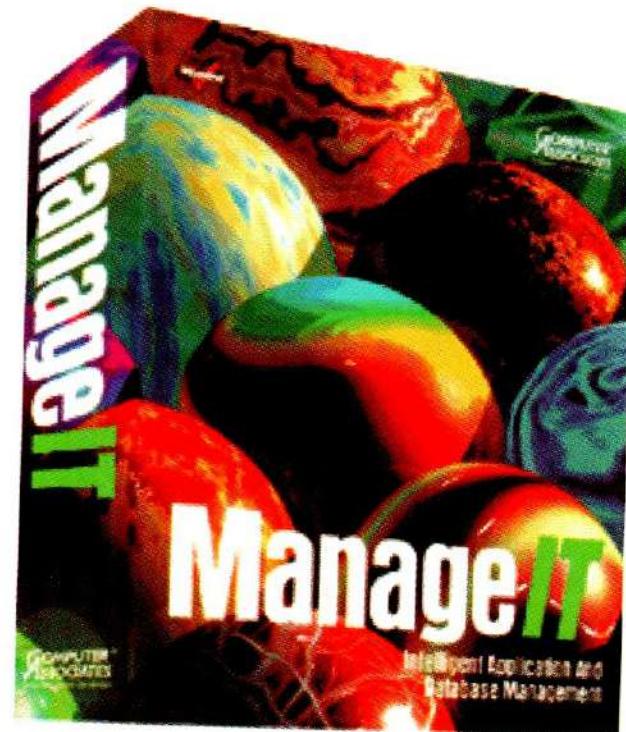
FYI No. 195

Computer Associates International Inc.

Computer Associates International Inc. (CA) released ManageIT, an application-management family. ManageIT detects anomalies in application or database behaviors before they reach end users and clients.

Using the Unicenter TNG Framework, ManageIT Knowledge Agents gather data on various components and processes within critical applications. Neugents ("neural agents" designed for predictive management), in turn, automatically and comprehensively analyze the data, instantly identifying patterns and anomalies that could lead to deterioration in performance or failure. As a result, clients can monitor application parameters to pinpoint the root

cause of problems and automatically take corrective action.



CA Associates Intl. Inc.'s ManageIT

CA provides Knowledge Agents for Enterprise Resource Planning applications, groupware, e-mail and middleware, as well as for underlying operating systems.

Contact Computer Associates International Inc., One Computer Associates Plaza, Islandia, NY 11749. Tel.: 516-342-5224 or 800-225-5224, e-mail: info@ca.com, Web: <http://www.ca.com>.

FYI No. 3

**When you
absolutely,
positively
gotta know.**

Read all about it.

Back content online.
<http://www.performance-computing.com>

PERFORMANCE
COMPUTING

Matthew Huff

FILE SMILES

Q. I have problems when I connect an NT 4.0 client to an HP-UX server using RAS through a null-modem cable. The NT client sends the keyword client and (I assume) is waiting for a response from HP-

UX. I have tried the same procedure between two NT 4.0 clients and it works properly. Is there any software or script in HP-UX that can make this connection possible?

A. Microsoft's KnowledgeBase article Q122318, using null-modem for RAS with third-party PPP Clients (<http://support.microsoft.com/support/kb/articles/q122/3/18.asp>), details the recommended changes to *modem.inf* that are required because of Microsoft's assumption that null modems will only be used to connect two NT systems. It also documents the requirement that you send \$ from the client to start the conversation.

Normally, I would strongly recommend using the open-source pppd (<ftp://cs.anu.edu.au/pub/software/ppp/>) because of the strong documentation, debugging, and feature set. Unfortunately, it is not available for HP-UX. You might want to try iij-ppp (available on the ppp for HP-UX home page at <http://www.interex.org/~borg/ppp.html>).

Q. We have an NT server loaded with IIS and a UNIX server that has APIs which can be called through shell scripts. Both the UNIX and NT server are connected through TCP/IP. My query is, how can I call these APIs or shell script through my NT server?

A. Since the API is available from UNIX shell scripts, you can use rsh.exe from the IIS server to the UNIX host. You will need to add the NT user and machine to the UNIX host's *host.equiv* file. The rsh protocol,

however, has limited security. I would recommend using Wietse's TCP wrapper tools (<ftp://ftp.porcupine.org/pub/security/index.html>). Also, the return code from rsh is whether the job has successfully been started, rather than the return code of the program. One trick is to redirect the output of the executable to a file, and use standard out for the return code:

```
C:\>rsh alamo -n "ls >ls.out; echo $?"  
0
```

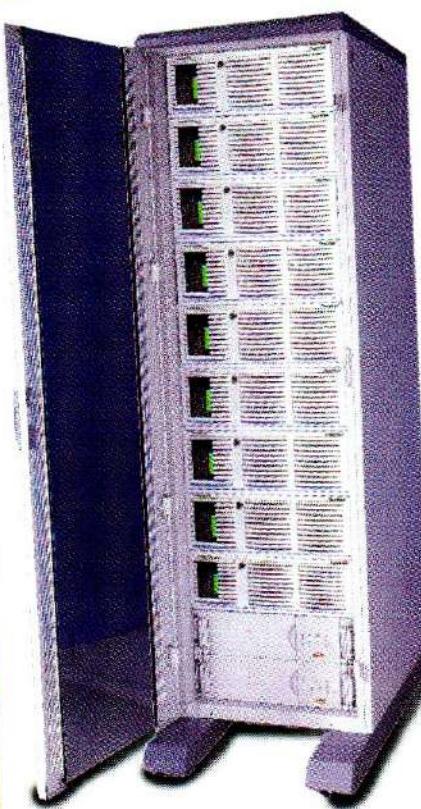
Q. How can I automatically fax documents in HP PCL5 format from UNIX?

A. If it were a question of faxing PostScript documents, a good solution would be to use the freeware ghostscript, netpbm, and HylaFAX (<http://www.hylafax.org>). Unfortunately, PCL doesn't have much support in UNIX as an input format. However, at least one commercial package, Faximum (<http://www.faximum.com>), does.

Q. I'm a novice at the UNIX end of an NT network. Although we pipe a UNIX machine folder to an NT server tape backup, the UNIX cron does a much more efficient and automatic job. I'm looking for a SCSI tape unit that will run on the UNIX box. What's the difference between UNIX SCSI and NT SCSI?

A. Most SCSI incompatibilities are related to different generations or features of SCSI technology, for example, SCSI-1 vs. SCSI-2, singled-ended vs. differential, and Wide vs. narrow, rather than NT vs. UNIX. Careful attention in matching the SCSI controller to the device prevents most problems. In most cases, NT SCSI devices will work fine on UNIX systems. In some rare cases, however, vendors have customized firmware to support special features of particular hardware that could cause difficulties. Most SCSI tape vendors, such as Exabyte, HP, and Quantum have tape technologies that are compatible with NT and UNIX systems.





Company:
Honeywell,
Consumer Products Division

Application:
Oracle Financial and
Manufacturing Applications

IT Team:
Kerry Asquith - Oracle Project Manager
Peggy Frederick - Director of Information Technology
Tim Wells - Network Manager

Server:
HP/UX K570



FlashDisk RAID:
400GB of FlashDisk RAID storage 4 dual host subsystems, each with eleven 9GB drives and 128MB cache

"FlashDisk eliminated Oracle's disk related I/O bottlenecks and saved us over \$250,000."

- Kerry Asquith

At Honeywell's Consumer Products Division, Oracle Financial and Manufacturing Applications performance during peak periods was simply unacceptable - and no amount of fine tuning seemed to provide a solution. "We were running Oracle Apps on DEC Alpha with DEC RAID. We added more processors, then we tried spreading our hot data and disks across controllers, rearranged table spaces, and spent months re-striping and re-mirroring," said Kerry Asquith, Oracle project manager. "Still, no breakthroughs."

Finally, Honeywell put a small Winchester Systems FlashDisk RAID demo to the test and moved

a quarter of their hot files to the FlashDisk. "With no other changes, response times were immediately 4 to 6 times faster," said Asquith. FlashDisk RAID uses high performance IBM UltraStar disks and massively parallel architecture to achieve lightning fast I/O response. Each FlashDisk delivers up to 9,799 read/writes per second - 10 times faster than a typical array. "Before, most transactions took up to 5 minutes on Oracle, now all transactions have dropped to under 5 seconds." In addition, FlashDisk also met Honeywell's need for fault tolerance - users can hot swap fans, power supplies, AC power and host cables.

FlashDisk® Case Study

Discover how FlashDisk can provide you with performance that no amount of application tuning can match - call Winchester Systems at **800.325.3700**. If you qualify, we'll conduct a **free system audit** to identify I/O bottlenecks and recommend solutions. Or visit our web site at www.winsys.com.

**APPROVED
FOR PURCHASE**
AUTHORIZED BY KA

WINCHESTER SYSTEMS®

400 West Cummings Park • Woburn, MA 01801

Predict The Future, And You Can Change The World. With NEUGENTS, You Can Do Both.



Introducing software that can emulate a human brain.

It's true. Neugents™ think like a human — only faster.

Neugents compute not in thousandths of a second, but millionths. Without emotion, subjectivity, or bias.

Neugents can analyze, make decisions, take action. They can process massive amounts of chaotic data and instantly identify complex patterns and relationships. Figuring out why things happen, and more importantly, predict what will happen next.

Neugents can learn. Using a unique self-learning algorithm, Neugents get smarter every second, every hour, every day.

The secret is Neural Network Technology. Computer Associates has taken this powerful technology and turned it into a patented application that goes far beyond traditional forecasting methods and rules-based applications.

Neugents can tell you what your sales are going to be next week, next month, or next year. They can tell you before your next product fails — and why. And for enterprise management, Neugents can do everything from warn you before a server goes down — or tell you when and where your next security breach will occur.

With virtually every aspect of your business, Neugents can accurately and consistently predict the future.

And when you can predict the future, you can not only change the future, you can change the world.

Call **1-877-Neugents** for more information.

NEUGENTS. Software That Can Think.™

**COMPUTER
ASSOCIATES**
Software superior by design.

©1998 Computer Associates International, Inc., Islandia, NY 11788-7000. All product names referenced herein are trademarks of their respective companies.

Select No. 3 at www.performancecomputing.com/fyi/